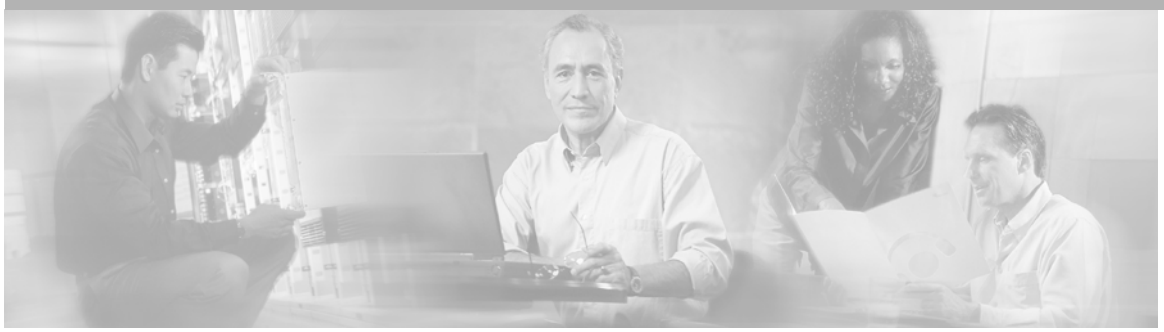

Enterprise Mobility 3.0 Design Guide

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number:
Text Part Number: OL-11573-01

Enterprise Mobility 3.0 Design Guide

Preface

Document Purpose

The purpose of this document is to describe the design and implementation of the Cisco Unified Wireless Network for the enterprise.

Intended Audience

This publication is for experienced network administrators who are responsible for design and implementation of wireless networks.

Document Organization

The following table lists and briefly describes the chapters of this guide.

Section	Description
Chapter 1, "Cisco Unified Wireless Network Solution Overview."	Summarizes the benefits and characteristics of the Cisco Unified Wireless Network for the enterprise.
Chapter 2, "Cisco Unified Wireless Technology and Architecture."	Discusses the key design and operational considerations in an enterprise Cisco Unified Wireless Deployment.
Chapter 3, "WLAN Radio Frequency Design Considerations."	Describes the basic radio frequency (RF) information necessary to understand RF considerations in various wireless local area network (WLAN) environments.
Chapter 4, "Cisco Unified Wireless Security."	Describes the natively available 802.11 security options and the advanced security features in the Cisco Unified Wireless solution, and how these can be combined to create an optimal WLAN solution.
Chapter 5, "Cisco Unified Wireless QoS."	Describes quality of service (QoS) in the context of WLAN implementations.

Section	Description
Chapter 6, "Cisco Unified Wireless Multicast Design."	Describes the improvements that have been made in IP multicast forwarding and provides information on how to deploy multicast in a wireless environment.
Chapter 7, "Cisco Unified Wireless Hybrid REAP."	Describes the Cisco Centralized WLAN architecture and its use of H-REAP.
Chapter 8, "Cisco Unified Wireless Control System."	Describes the Cisco Wireless Control System (WCS) and addresses management considerations to consider when using it to design, deploy, and manage your enterprise wireless LAN.
Chapter 9, "Cisco Unified Wireless Security Integration."	Discusses the integration of wired network security into the Cisco Unified Wireless Solution.
Chapter 10, "Cisco Wireless Mesh Networking."	Describes the use of wireless mesh.
Chapter 11, "VoWLAN Design Recommendations."	Provide design considerations when deploying voice over WLAN (VoWLAN) solutions.
Chapter 12, "Cisco Unified Wireless Guest Access Services."	Describes the use of guest access services in the centralized WLAN architecture.
Chapter 13, "Mobile Access Router, Universal Bridge Client, and Cisco Unified Wireless."	Describes the use of the mobile access router, universal bridge client, and mesh networks.
Chapter 14, "Cisco Unified Wireless and Mobile IP."	Describes the inter-workings of the Cisco Mobile Client (CMC) over a Cisco Unified Wireless Network (WiSM).
Chapter 15, "Cisco Unified Wireless Location-Based Services."	Discusses the Cisco Location-Based Service (LBS) solution and the areas that merit special consideration involving design, configuration, installation, and deployment.

Modification History

Revision	Date	Originator	Comments

Cisco Unified Wireless Network Solution Overview

This chapter summarizes the benefits and characteristics of the Cisco Unified Wireless Network for the enterprise.

WLAN Introduction

The mobile user requires the same accessibility, security, quality of service (QoS), and high availability currently enjoyed by wired users. Whether you are at work, at home, on the road, locally or internationally, there is a need to connect. The technological challenges are apparent, but to this end, mobility plays a role for everyone. Companies are deriving business value from mobile and wireless solutions. What was once a vertical market technology is now mainstream, and is an essential tool in getting access to voice, real-time information, and critical applications such as e-mail and calendar, enterprise databases, supply chain management, sales force automation, and customer relationship management.

WLAN Solution Benefits

WLANs provide the user with a new way to communicate while accommodating the way business is done now. The benefits achieved by WLANs are the following:

- *Mobility within building or campus*—Facilitates implementation of applications that require an always-on network and that tend to involve movement within a campus environment.
- *Convenience*—Simplifies networking of large, open people areas.
- *Flexibility*—Allows work to be done at the most appropriate or convenient place rather than where a cable drop terminates. Getting the work done is what is important, not where you are.
- *Easier to set-up temporary spaces*—Promotes quick network setup of meeting rooms, war rooms, or brainstorming rooms tailored to variations in the number of participants.
- *Lower cabling costs*—Reduces the requirement for contingency cable plant installation because the WLAN can be employed to fill the gaps.
- *Easier adds, moves, and changes and lower support and maintenance costs*—Temporary networks become much easier to set up, easing migration issues and costly last-minute fixes.
- *Improved efficiency*—Studies show WLAN users are connected to the network 15 percent longer per day than hard-wired users.

- *Productivity gains*—Promotes easier access to network connectivity, resulting in better use of business productivity tools. Productivity studies show a 22 percent increase for WLAN users.
- *Easier to collaborate*—Facilitates access to collaboration tools from any location, such as meeting rooms; files can be shared on the spot and requests for information handled immediately.
- *More efficient use of office space*—Allows greater flexibility for accommodating groups, such as large team meetings.
- *Reduced errors*—Data can be directly entered into systems as it is being collected, rather than when network access is available.
- *Improved efficiency, performance, and security for enterprise partners and guests*—Promoted by implementing guest access networks.
- *Improved business resilience*—Increased mobility of the workforce allows rapid redeployment to other locations with WLANs.

Requirements of WLAN Systems

WLAN systems run either as an adjunct to the existing wired enterprise network or as a free-standing network within a campus or branch, individual tele-worker, or tied to applications in the retail, manufacturing, or health care industries. WLANs must permit secure, encrypted, authorized communication with access to data, communication, and business services as if connected to the resources by wire.

WLANs must be able to do the following:

- *Maintain accessibility to resources while employees are not wired to the network*—This accessibility enables employees to respond more quickly to business needs regardless of whether they are meeting in a conference room with a customer, at lunch with coworkers in the company cafeteria, or collaborating with a teammate in the next building.
- *Secure the enterprise from unauthorized, unsecured, or "rogue" WLAN access points*—IT managers must be able to easily and automatically detect and locate rogue access points and the switch ports to which they are connected, active participation of both access points, and client devices that are providing continuous scanning and monitoring of the RF environment.
- *Extend the full benefits of integrated network services to nomadic users*—IP telephony and IP video-conferencing are supported over the WLAN using QoS, which by giving preferential treatment to real-time traffic, helps ensure that the video and audio information arrives on time. Firewall and Intruder Detection that are part of the enterprise framework are extended to the wireless user.
- *Segment authorized users and block unauthorized users*—Services of the wireless network can be safely extended to guests and vendors. The WLAN must be able to configure support for a separate public network—a guest network.
- *Provide easy, secure network access to visiting employees from other sites*—There is no need to search for an empty cubicle or an available Ethernet port. Users should securely access the network from any WLAN location. Employees are authenticated through IEEE 802.1x and Extensible Authentication Protocol (EAP), and all information sent and received on the WLAN is encrypted.
- *Easily manage central or remote access points*—Network managers must be able to easily deploy, operate, and manage hundreds to thousands of access points within the WLAN campus deployments and branch offices or retail, manufacturing, and health care locations. The desired result is one

framework that provides medium-sized to large organizations the same level of security, scalability, reliability, ease of deployment, and management that they have come to expect from their wired LANs.

Wireless LANs in the enterprise have emerged as one of the most effective means for connecting to a network. Figure 1-1 shows the elements of the Cisco Unified Wireless Network.

Figure 1-1 Cisco Unified Wireless Network Architecture in the Enterprise

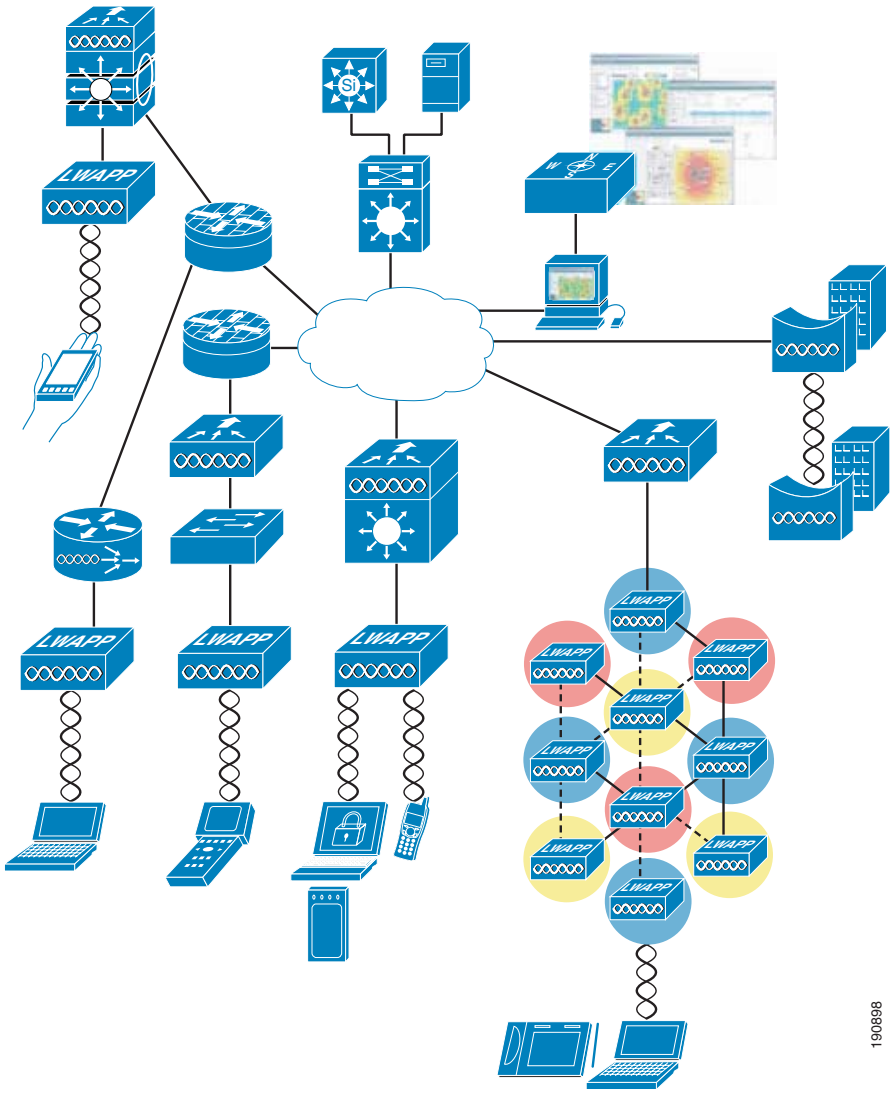

The following five interconnected elements work together to deliver a unified enterprise-class wireless solution:

- Client devices
- Access points
- Network unification
- World-class network management
- Mobility services

Beginning with a base of client devices, each element adds capabilities as network needs evolve and grow, interconnecting with the elements above and below it to create a comprehensive, secure WLAN solution.

The Cisco Unified Wireless Network cost-effectively addresses the wireless LAN (WLAN) security, deployment, management, and control issues facing enterprises. This framework integrates and extends wired and wireless networks to deliver scalable, manageable, and secure WLANs with the lowest total cost of ownership. The Cisco Unified Wireless Network provides the same level of security, scalability, reliability, ease of deployment, and management for wireless LANs that organizations expect from their wired LANs.

The Cisco Unified Wireless Network includes two secure, enterprise-class WLAN solutions. Customers can choose to deploy either Autonomous Cisco Aironet Access Points running Cisco IOS Software or Lightweight Access Points using a Cisco Wireless LAN Controller (WLC). The primary difference between these two types of access points lies in their implementation of access point control and management.

The devices are available in two versions: those configured for lightweight operation in conjunction with Cisco Wireless LAN Controllers and the Wireless Control System (WCS) as well as those configured for autonomous operation, used independently or in conjunction with the CiscoWorks Wireless LAN Solution Engine (WLSE). Autonomous access points along with the CiscoWorks WLSE deliver a core set of features. Autonomous access points may be field upgraded to lightweight operation and an advanced feature set. Customers can choose the access point that best meets their WLAN deployment needs today knowing that Cisco provides the investment protection and a migration path to evolve their WLAN going forward.

For more information about the Cisco Unified Wireless Network, see the following URL: http://www.cisco.com/go/unifiedwireless

Cisco Unified Wireless Network

The core feature set includes autonomous Cisco Aironet access points, the Wireless Control System (WCS), and Wireless LAN Controllers (WLC), including the Cisco Catalyst 6500 Wireless Services Module (WiSM), the 440X, and 2006 controls, the WLCM ISR module, and the WS-C3750G integrated controller.

The core feature set is deployable in the following configurations today:

- APs and WLC
- APs, WLCs, and WCS
- APs, WLC, WCS, and LBS

Adding optional Cisco Compatible Extensions client devices provides additional benefits, including advanced enterprise-class security, extended RF management, and enhanced interoperability.

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/products/ps6366/prod_technical_reference09186a00806cfa96.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

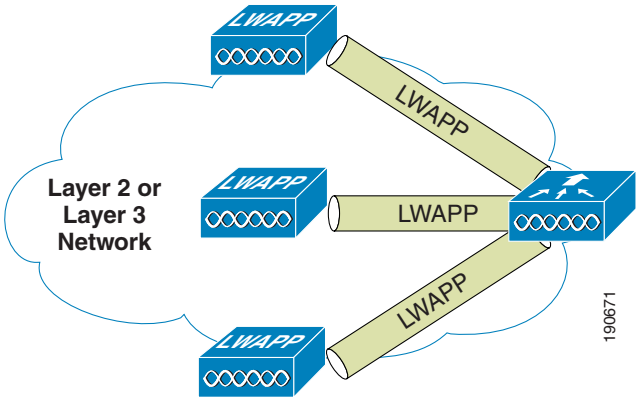

Although the LWAPP protocol has a number of components, only the components of the LWAPP protocol that impact the network design and operation are discussed in this document.

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

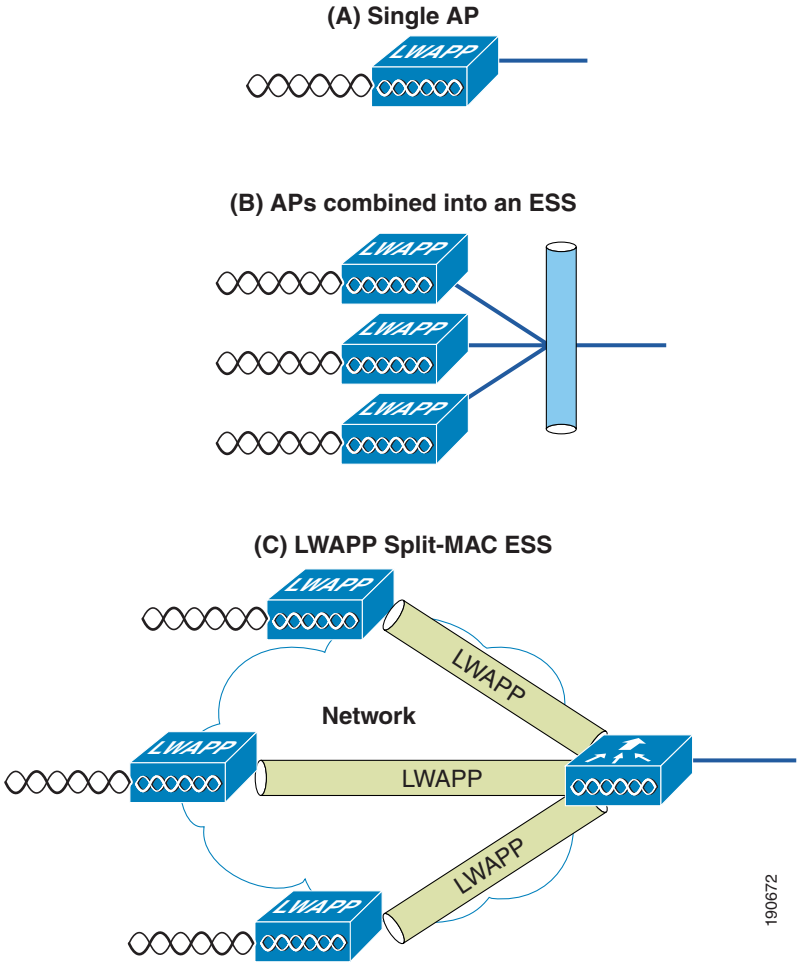

The simple timing-dependent operations are generally managed on the LWAPP AP, and more complex and less time-dependent operations are managed on the WLC.

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

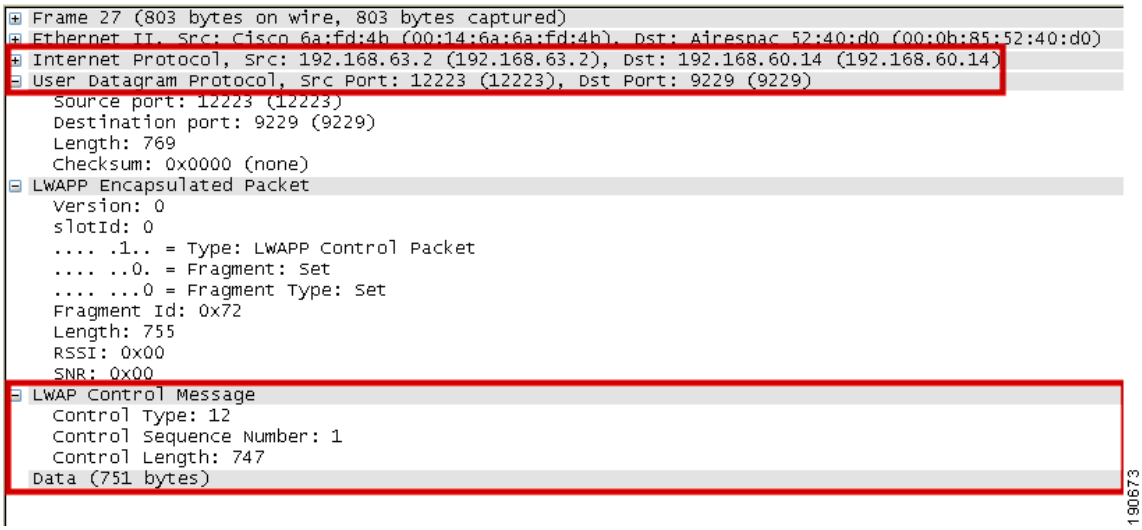

Figure 2-4 shows a decode of an LWAPP packet containing an 802.11 probe request. This packet is from the LWAPP AP to the WLC, and uses UDP port 12222, as do all LWAPP-encapsulated 802.11 frames. In this case, RSSI and SNR values are also included in the LWAPP packet to provide RF information to the WLC.

Figure 2-4 802.11 Probe Request in LWAPP

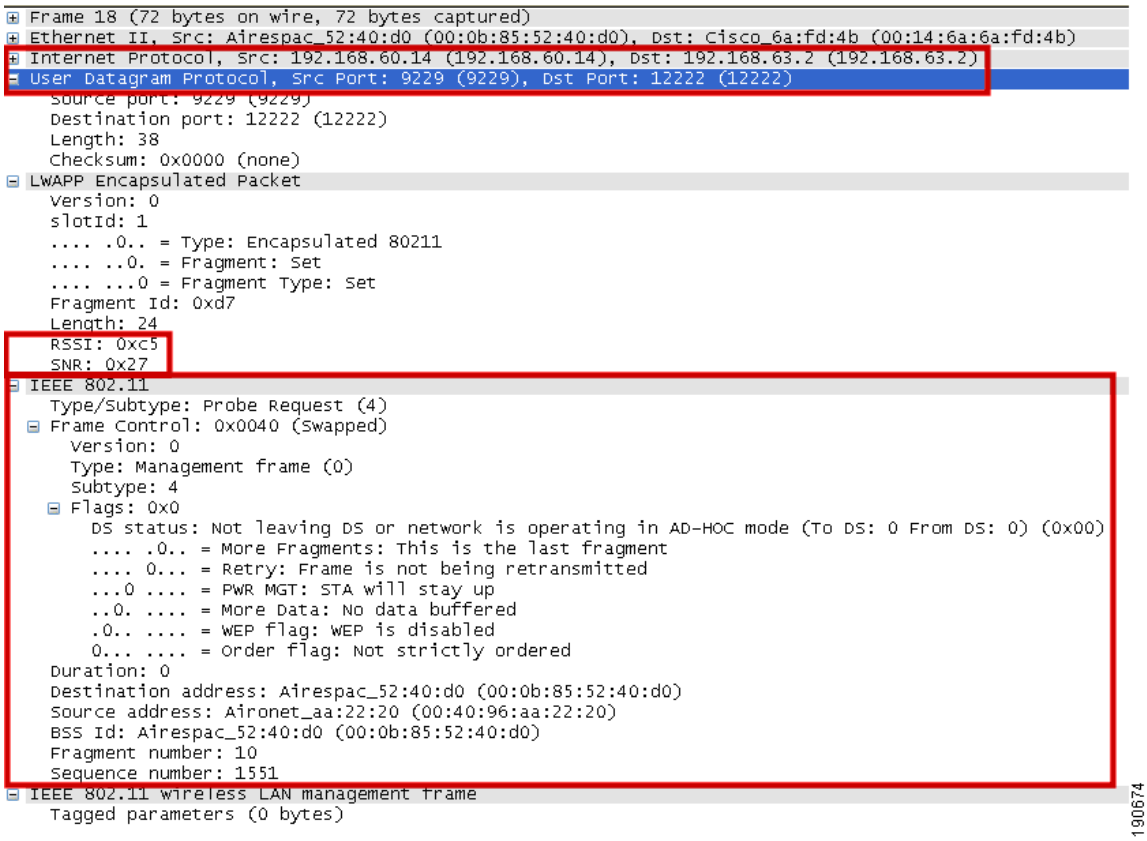

Figure 2-5 shows another LWAPP-encapsulated 802.11 frame, but in this case it is an 802.11 data frame, like that shown in Figure 2-4. It contains the complete 802.11 frame, as well as the RSSI and SNR information for the WLC, and is primarily shown here to demonstrate that the 802.11 data frame is treated 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

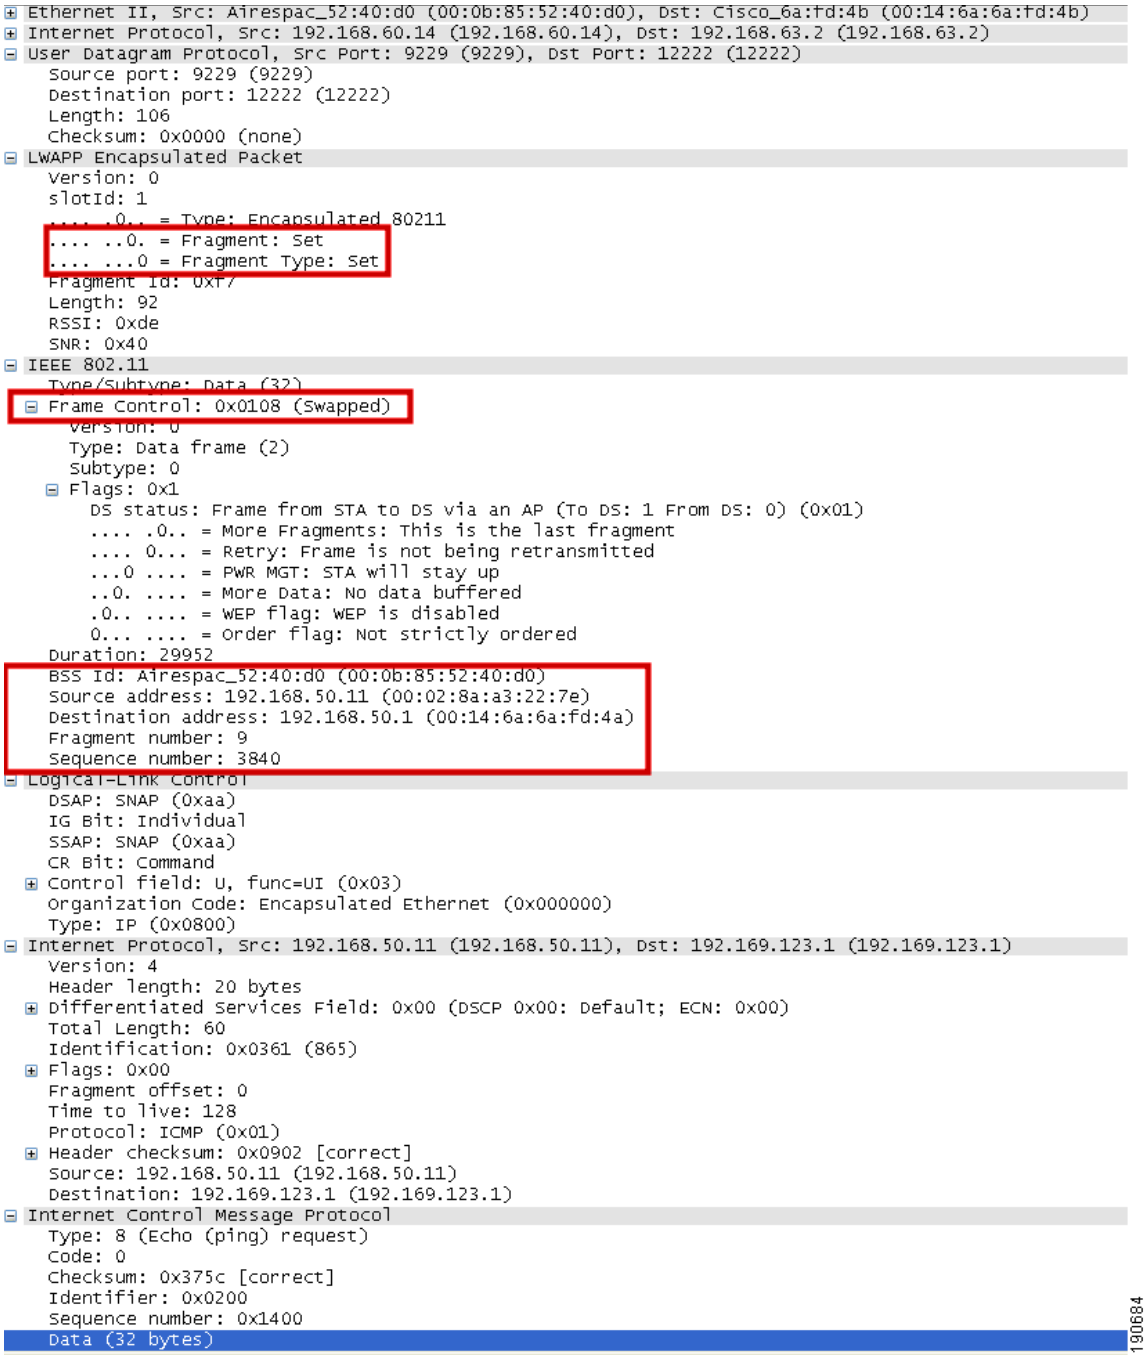

WLC Discovery and Selection

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/products/ps6366/prod_technical_reference09186a00806cfa96.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	
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****

Table 2-3 AP Comparison (2)

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/partner/products/ps6108/prod_brochure0900aecd8035a015.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/products/ps6366/prod_technical_reference09186a00806cfa96.html
- Cisco Wireless LAN Controller Configuration Guide, Release 4.0—
http://www.cisco.com/en/US/products/ps6366/products_configuration_guide_book09186a00806b0077.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

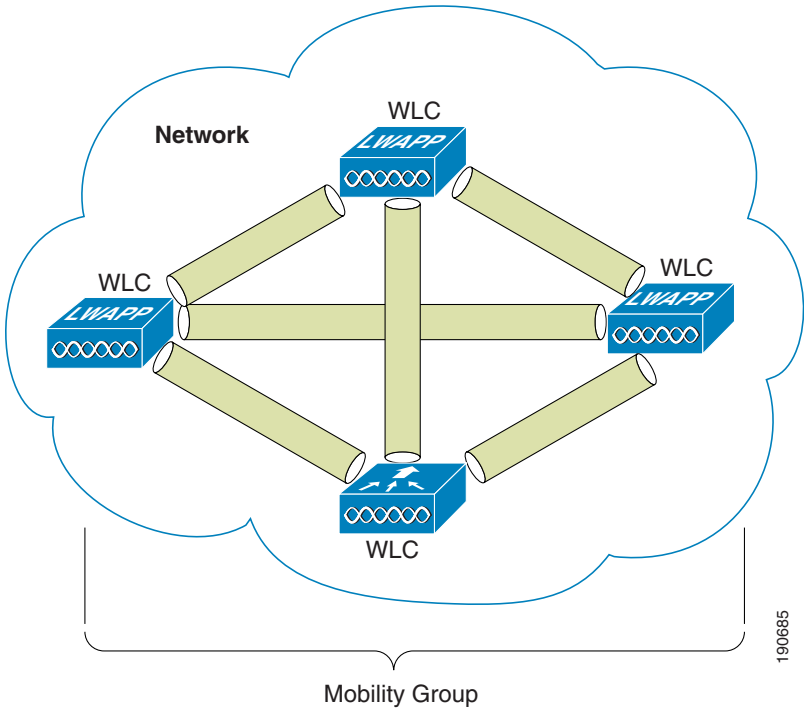

Creating Mobility Group

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

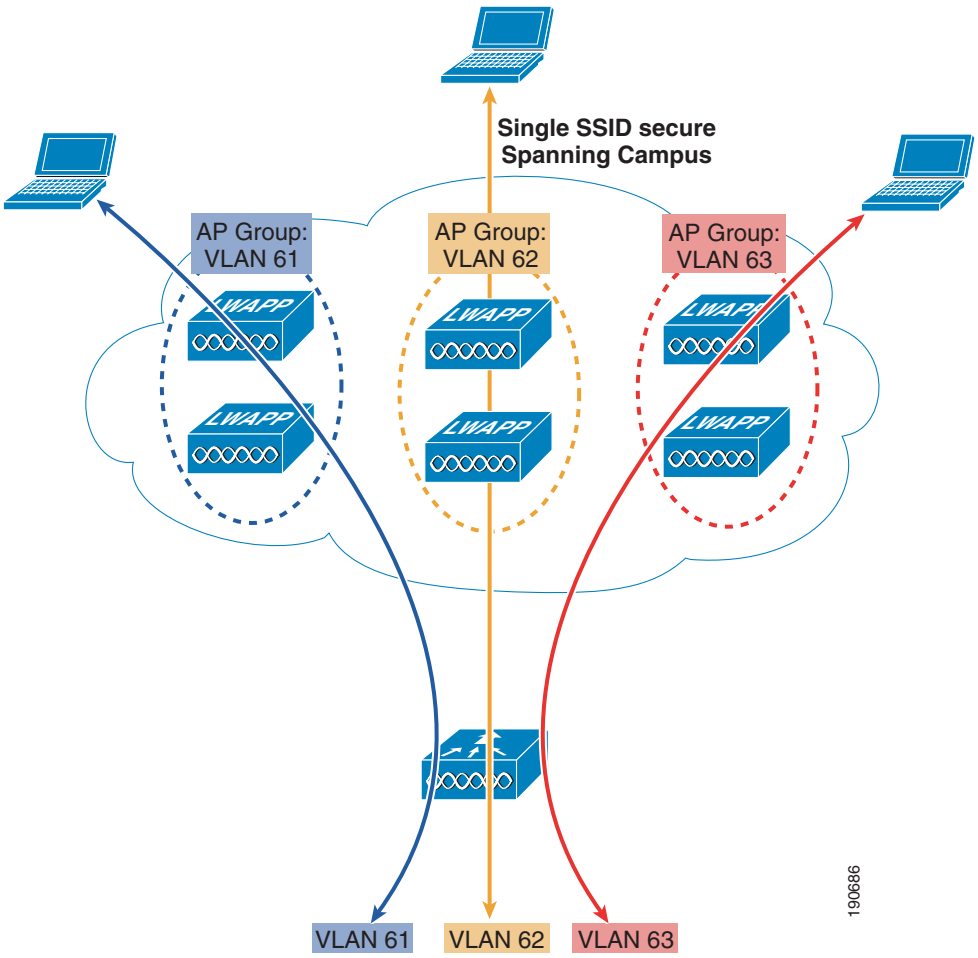

In Figure 2-7, there are three dynamic interfaces configured, mapping to three site-specific VLANs: VLANs 61, 62, and 63. These site-specific VLANs apply to the secure SSID for normal corporate users. A corporate user associating to the secure SSID on an AP in the AP Group corresponding to VLAN 61 gets 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

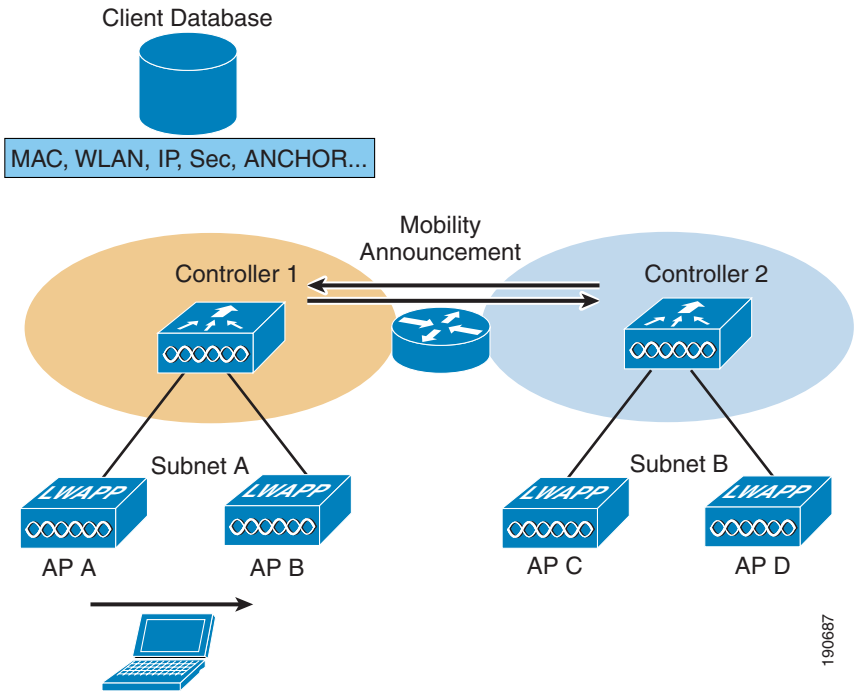

In instances when there are multiple WLCs connected to the same subnet, and therefore a client can roam between WLCs but remain on the same subnet, mobility announcements are passed between the related WLCs to pass client context information between WLCs. This WLC then becomes the anchor WLC for that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

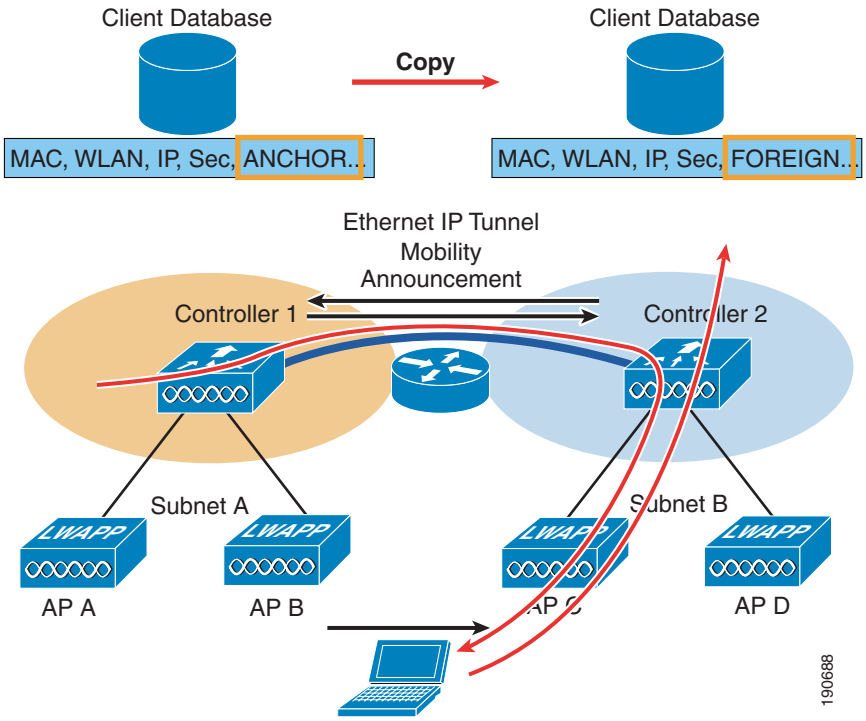

The client begins with a connection to AP B on WLC 1. This creates an ANCHOR entry in the WLC client database. As the client moves away from AP B and makes an association with AP C, WLC 2 sends a mobility announcement to peers in the mobility group looking for the WLC with the client MAC address. WLC 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

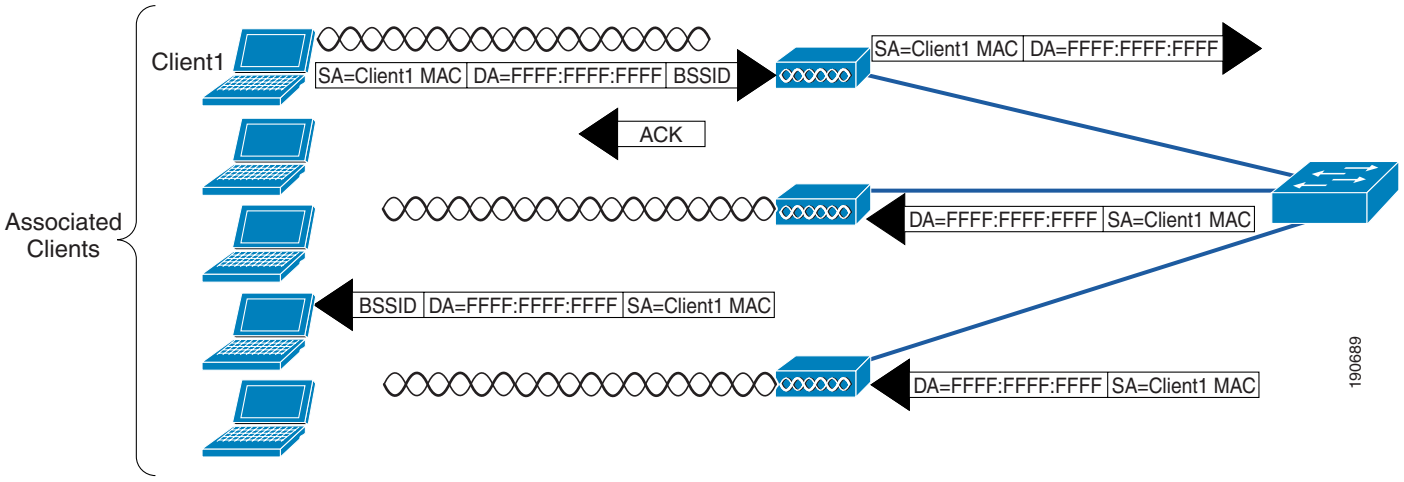

If there are other APs on the same wired VLAN as the AP of Figure 2-10, they forward the wired broadcast packet out their wireless interface.

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

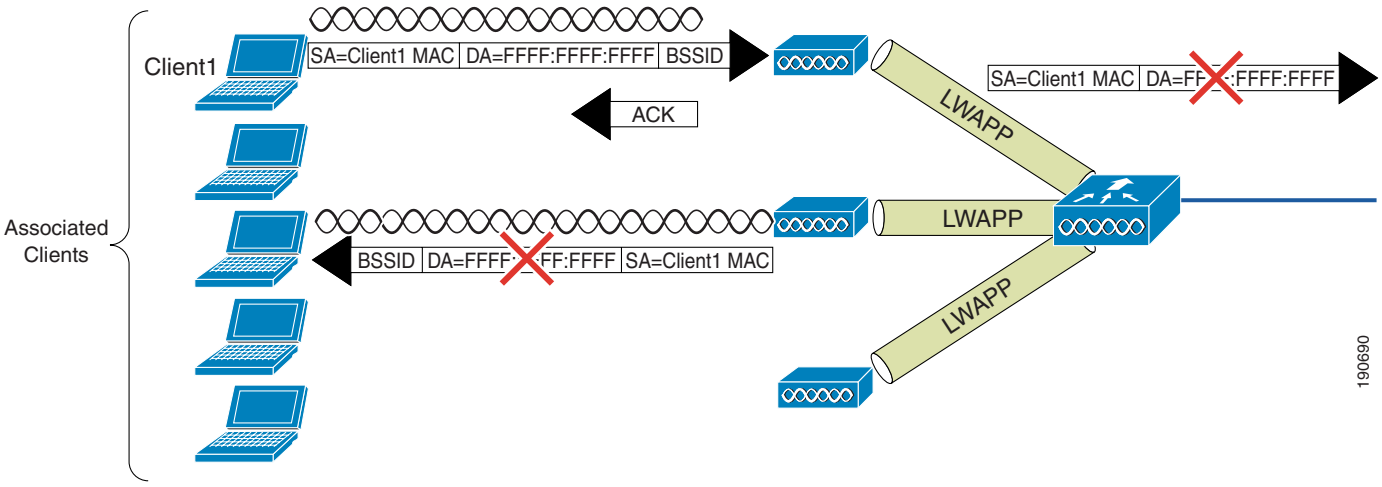

WLC Broadcast and Multicast Details

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

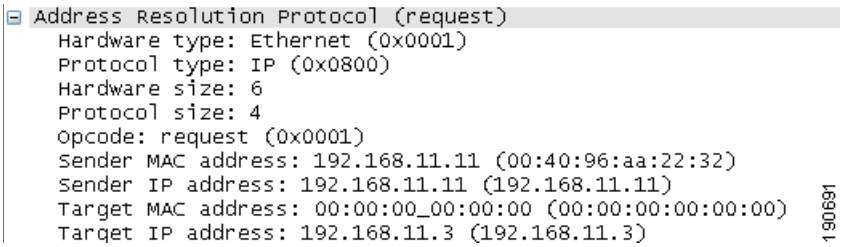

On seeing the ARP request, the WLC either responds directly, acting as an ARP proxy, or forwards the request out the wired interface to have it resolved by another WLC; the WLC does not forward the ARP broadcast back out to the WLAN.

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 Consideration

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

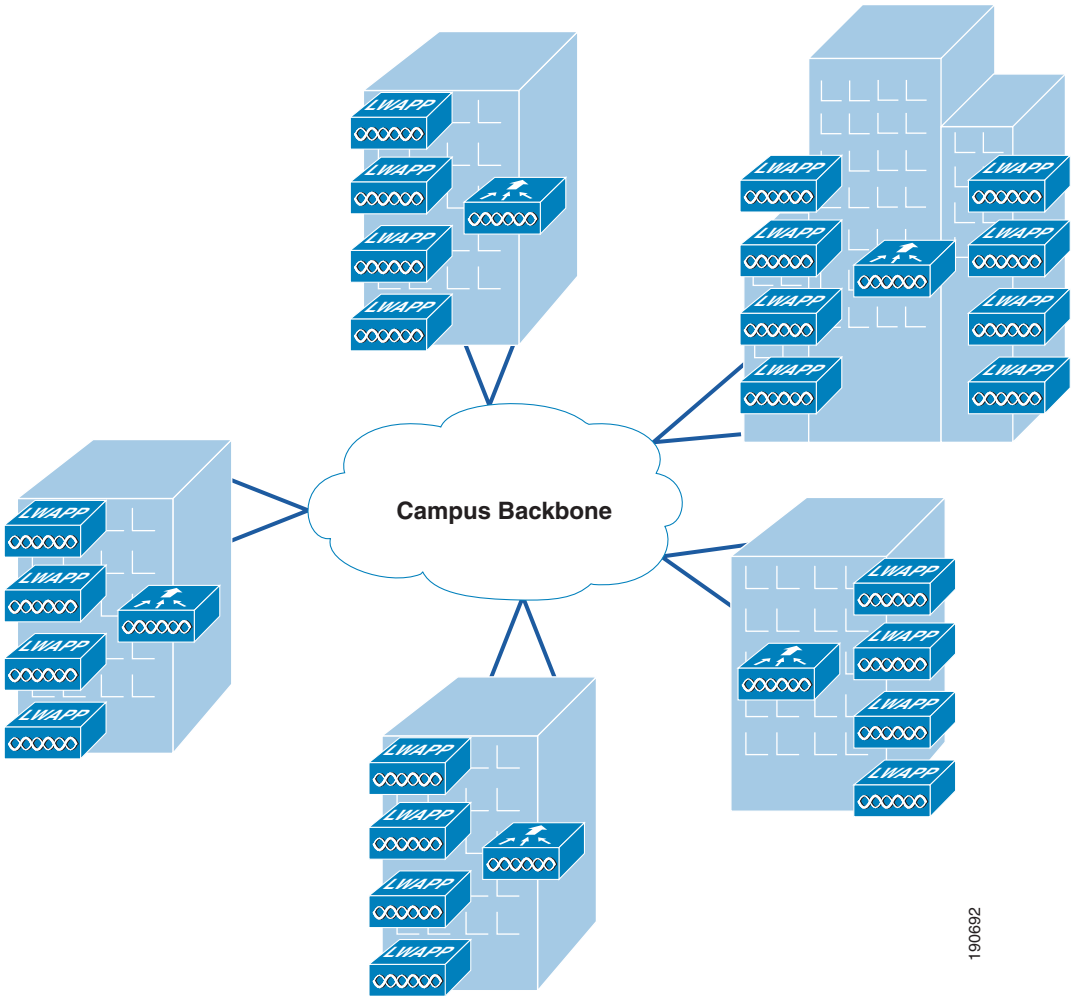

- Centralized WLC deployment—WLCs are placed in a centralized location in the network where most LWAPP tunnels between APs and WLCs must traverse the campus backbone network. A schematic of a centralized WLC deployment is shown in Figure 2-14. Note that the centralized WLC (a pair of WiSMs enabled 6500s 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

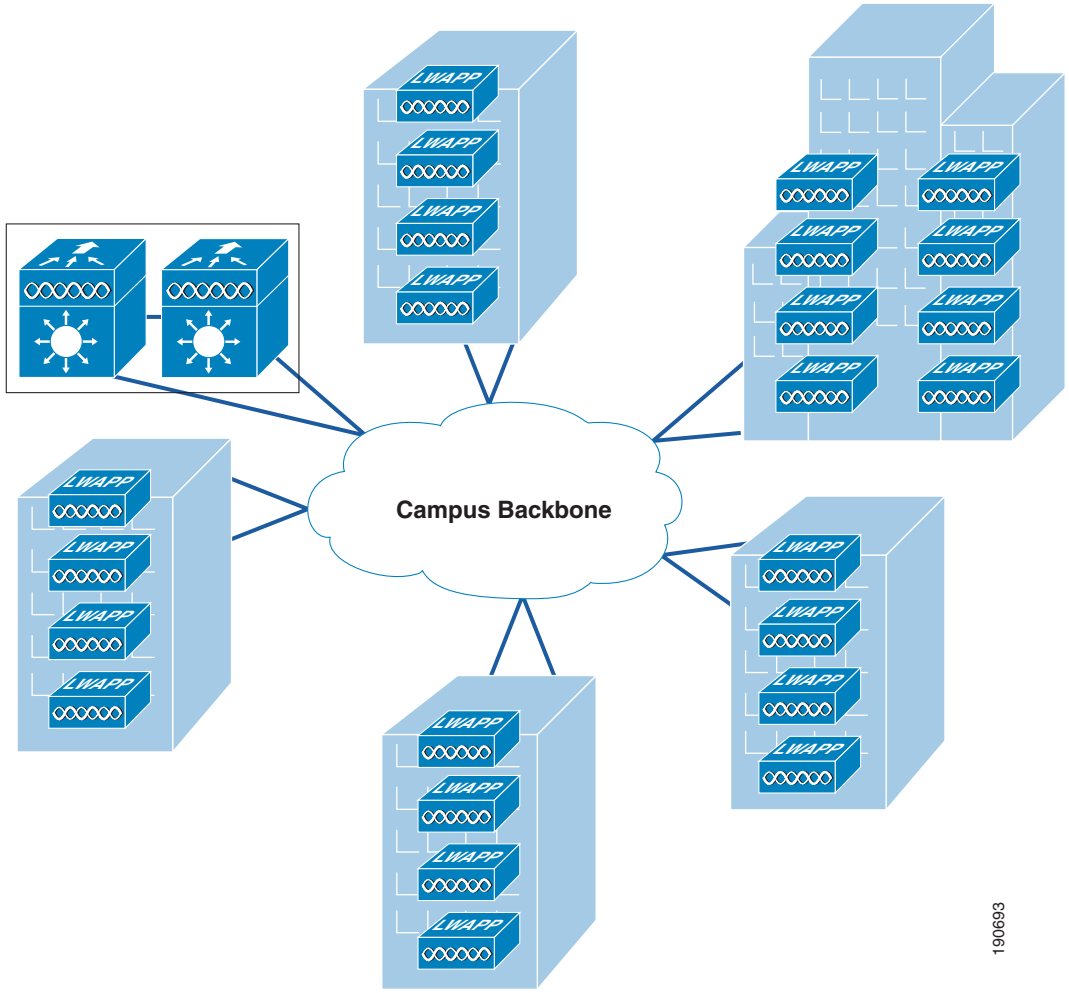

Centralizing WLCs

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

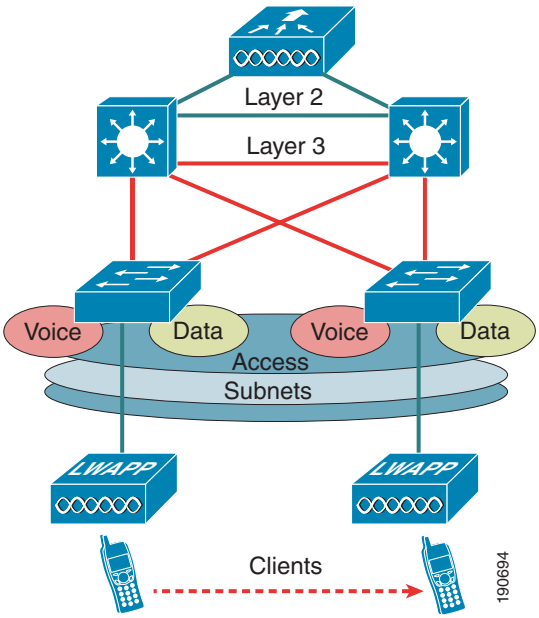

This recommendation is made for a number of reasons, including the following:

- 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

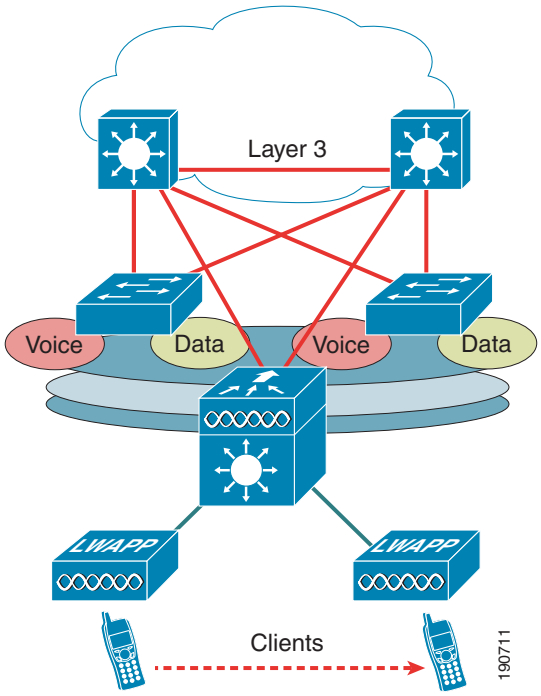

Link Budget and Wired Network Performance

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.

WLAN Radio Frequency Design Considerations

Introduction

This chapter describes the basic radio frequency (RF) information necessary to understand RF considerations in various wireless local area network (WLAN) environments. This chapter includes information on the following topics:

- Regulatory domains and frequencies
- Understanding the IEEE 802.11 standards
- RF spectrum implementations including 802.11b/g and 802.11a
- Planning for RF deployment
- Manually fine-tuning WLAN coverage
- Radio Resource Management (RRM), also known as Auto-RF

RF Basics

This section provides a summary of regulatory domains and their operating frequencies.

Regulatory Domains

Devices that operate in unlicensed bands do not require any formal licensing process, but when operating in these bands, the user is obligated to follow the government regulations for that region. The regulatory domains in different parts of the world monitor these bands according to different criteria, and the WLAN devices used in these domains must comply with the specifications of the relevant governing regulatory domain. Although the regulatory requirements do not affect the interoperability of IEEE 802.11b/g and 802.11a-compliant products, the regulatory agencies do set certain criteria in the standard. For example, the emission requirements for WLAN to minimize the amount of interference a radio can generate or receive from another radio in the same proximity. It is the responsibility of the vendor to get the product certified from the relevant regulatory body. Table 3-1 summarizes the current regulatory domains for Wi-Fi products. The main regulatory domains are FCC, ETSI, and the MKK.

Besides following the requirements of the regulatory agencies, many vendors also ensure compatibility with other vendors through the Wi-Fi certification program (www.wi-fi.org).

For a complete listing of other countries' regulatory transmit power settings and allowed frequency use, see the following URL:

http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_chapter09186a008059c96f.html

Table 3-1 Regulatory Domains

Regulatory Domain	Geographic Area
Americas or FCC (United States Federal Communication Commission)	North, South, and Central America, Australia and New Zealand, various parts of Asia and Oceania
Europe or ETSI (European Telecommunications Standards Institute)	Europe (both EU and non EU countries), Middle East, Africa, various parts of Asia and Oceania
Japan (MKK)	Japan
China	People's Republic of China (Mainland China)
Israel	Israel
Singapore[1]	Singapore
Taiwan[1]	Republic of China (Taiwan)

[1] The regulations of Singapore and Taiwan for wireless LANs are particular to these countries only for operation in the 5 GHz band. Singapore and Taiwan are therefore only regulatory domains for 5 GHz operation; for operation in 2.4 GHz, they fall into the ETSI and FCC domains, respectively.

Note See the Cisco website for compliance information and also check with your local regulatory authority to find out what is permitted within your country. The information provided in Table 3-2 and Table 3-3 should be used as a general guideline. For up-to-date information on which Cisco products meet regional requirements, see the following URL:
http://www.cisco.com/warp/public/779/smbiz/wireless/approvals.html#4

Operating Frequencies

The 802.11b/g band regulations have been relatively constant, given the length of time it has been operating. The FCC allows for 11 channels, ETSI allows for up to 13 channels, and Japan allows up to 14 channels, but requires a special license to operate in channel 14.

For 802.11a, countries are moving to open the frequency range 5.250–5.350 GHz (UNII-2) and the frequency range 5.470 to 5.780 GHz for additional 802.11a channels. These various frequencies are covered in more detail in the specific 802.11 sections in this chapter.

802.11b/g Operating Frequencies and Data Rates

Ratified in September 1999, the 802.11b standard operates in the 2.4 GHz spectrum and supports data rates of 1, 2, 5.5, and 11 Mbps. 802.11b enjoys broad user acceptance and vendor support. 802.11b technology has been deployed by thousands of enterprise organizations, which typically find its speed and performance acceptable for their current applications.

The 802.11g standard, which was ratified in June 2003, operates in the same spectrum as 802.11b and is backward-compatible with the 802.11b standard. 802.11g supports the additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. 802.11g delivers the same 54 Mbps maximum data rate as 802.11a, but operates in the same 2.4 GHz band as 802.11b. It also provides backward compatibility with existing 802.11b devices.

Table 3-2 lists the various 802.11b/g channel frequencies and specifies whether a regulatory agency allows their use in their domain. Note that not all of these frequencies are available for use in all regulatory domains.

Table 3-2 Operating Frequency Range for 802.11b and 802.11g

Channel Identifier	Center Frequency	FCC (America)	ESTI (EMEA)	TELEC (Japan)	MOC (Israel Outdoor)[1]
1	2412	X	X	X	
2	2417	X	X	X	
3	2422	X	X	X	
4	2427	X	X	X	
5	2432	X	X	X	X
6	2437	X	X	X	X
7	2442	X	X	X	X
8	2447	X	X	X	X
9	2452	X	X	X	X
10	2457	X	X	X	X
11	2462	X	X	X	X
12	2467		X	X	X
13	2472		X	X	X
14[2]	2484			X	

[1] Israel allows channels 1 through 13 indoors.

[2] Japan requires a special license for channel 14.

802.11a Operating Frequencies and Data Rates

Operating in the unlicensed portion of the 5 GHz radio band, 802.11a is immune to interference from devices that operate in the 2.4 GHz band, such as microwave ovens, cordless phones, and Bluetooth (a short-range, low-speed, point-to-point, personal-area-network wireless standard). Because the 802.11a standard operates in a different frequency range, it is not compatible with existing 802.11b or 802.11g-compliant wireless devices, but it does mean that 2.4-GHz and 5-GHz equipment can operate in the same physical environment without interference.

Choosing between these two technologies (802.11b/g and 802.11a) does not involve a one-for-one trade-off. They are complementary technologies and will continue to coexist in future enterprise environments. Those responsible for implementing these technologies must be able to make an educated choice between deploying 2.4 GHz-only networks, 5 GHz-only networks, or a combination of both. Organizations with existing 802.11b networks cannot simply deploy a new 802.11a network for existing

APs and expect to have their 802.11a 54 Mbps coverage in the same areas as their 11Mbps 802.11b coverage. The technical characteristics of both these bands simply do not allow for this kind of coverage interchangeability.

802.11a provides data rates of 6, 9, 12, 18, 24, 36, 48, with a maximum data rate of 54 Mbps, though generally at shorter ranges for a given power and gain, but it has up to 23 nonoverlapping frequency channels (depending on the geographic area) compared to the three nonoverlapping channels of 802.11b/g, which results in increased network capacity, improved scalability, and the ability to create microcellular deployments without interference from adjacent cells.

The 5 GHz band in which 802.11a operates is divided into several different sections. Each of the Unlicensed National Information Infrastructure (UNII) bands presented in Table 3-3 was originally intended for different uses, but all can currently be used by indoor 802.11a with appropriate power restrictions. Initially, the FCC defined only the UNII-1, UNII-2, and UNII-3 bands, each of which had four channels. The channels were spaced 20 MHz apart with an RF spectrum bandwidth of 20 MHz, thereby providing nonoverlapping channels.

There are differing limitations on these three UNII bands. Restrictions vary between them for transmit power, antenna gain, antenna styles, and usage. The UNII-1 band is designated for indoor operations, and initially had a restriction of permanently attached antennas. The UNII-2 band was designated for indoor or outdoor operations, and permitted external antennas**.** The UNII-3 band was intended for outdoor bridge products and permitted external antennas, but the UNII-3 band can now be used for indoor or outdoor 802.11a WLANs as well.

The channels in UNII-1 (5.150 to 5.250 GHz) are 34, 36, 38, 40, 42, 44, 46, and 48. The channels in UNII-2 (5.250–5.350 GHz) are 52, 56, 60, 64 and require Dynamic Frequency Selection (DFS) and Transmitter Power Control (TPC). The channels in the new frequency range (5.470–5.725 GHz) are 100, 104, 108, 112, 116, 120, 124, 128, 132, 136, 140 and require DFS and TPC. The channels in UNII-3 are 149, 153, 157, 161, 165 and require DFS and TPC. Not all channels in a given range can be used in all of the regulatory domains. Figure 3-1 shows the various channels in the UNII-1, 2, and 3 bands, along with the additional 11 new channels.

For more information on FCC regulation updates, see the following URL:

http://www.cisco.com/en/US/products/hw/wireless/ps469/products_white_paper0900aecd801c4a88.shtml

Table 3-3 Operating Frequency Range for 802.11a

Channel Identifier	36	40	44	48	52	56	60	64	149	153	157	161
Center Frequency	5180	5200	5220	5240	5260	5280	5300	5320	5745	5765	5785	5805
Band	**UNII-1**				**UNII-2**				**UNII-3**			

Table 3-3 shows the standard 802.11a frequencies. Table 3-4 shows the specific frequency bands and channel numbers for a few specific regulatory domains.

Table 3-4 Additional Frequency Bands and Channel Numbers for Other Regulatory Domains

Regulatory Domain	**Frequency Band**	**Channel Number**	**Center Frequency**
Japan[1]	U-NII lower bands	36	5.180
		40	5.200
		44	5.220
		48	5.240

Table 3-4 Additional Frequency Bands and Channel Numbers for Other Regulatory Domains

Singapore	U-NII lower band	36	5.180
		40	5.200
		44	5.220
		48	5.240
Taiwan		52	5260
		56	5280
		60	5300
		64	5320
EMEA 1 Australia New Zealand	Same as USA	Same as USA	Same as USA
EMEA 2[2]	U-NII lower band	36	5.180
		40	5.200
		44	5.220

[1] Japan is changing from channels 34,38,42,46 to 36,40,44,48 and adding channels 52, 56, 60, and 64 but required Dynamic Frequency Selection on channels 52 - 64. Cisco Equipment with the old channels are designated with a -J, Cisco equipment with the new channels are designated with a -P.

[2] Some EMEA countries, such as Denmark and Germany, are limited to 20 mW.

In February of 2004, the FCC released a revision to the regulations covering the 5 GHz 802.11a channel usage. This revision added 11 additional channels, bringing the available channels capacity to 23 channels (see Figure 3-1). The new additional 11 channels will be for Indoor/Outdoor use. To use the 11 new channels, however, radios must comply with two features that are part of the 802.11h specification: transmit power control (TPC) and dynamic frequency selection (DFS). DFS is required to avoid radar that operates in this frequency range, but it can also be used for other purposes, such as dynamic frequency planning. 802.11h has been supported since Cisco Unified Wireless Network Software Release 3.1.

For a complete listing of other countries regulatory transmit power settings and allowed frequency use, see the following URL:

http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_chapter09186a008059c96f.html

Figure 3-1 802.11 Channel Capacity

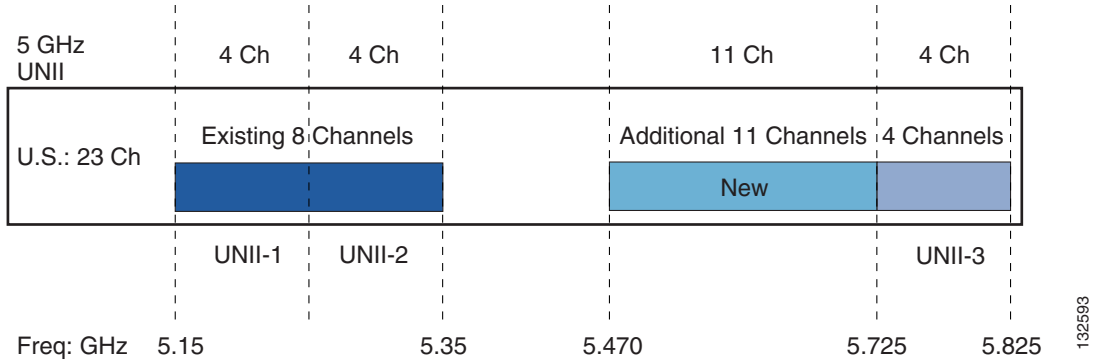

Understanding the IEEE 802.11 Standards

IEEE 802.11 is the working group within the Institute for Electrical and Electronics Engineers (IEEE) responsible for wireless LAN standards at the physical and link layer (Layer 1 and Layer 2) of the OSI model, as compared to the Internet Engineering Task Force (IETF, which works on network layer (Layer 3) protocols. Within the 802.11 working group are a number of task groups that are responsible for elements of the 802.11 WLAN standard. Table 3-5 summarizes some of the task group initiatives.

For more information on these working groups, see the following URL:

http://www.ieee802.org/11/

Table 3-5 IEEE 802.11 Task Group Activities

Task Group	Project
MAC	To develop one common MAC for WLANs in conjunction with a physical layer entity (PHY) task group
PHY	To develop three WLAN PHYs—Infrared, 2.4 GHz FHSS, 2.4 GHz DSSS
a	To develop PHY for 5 GHz UNII band
b	To develop higher rate PHY in 2.4 GHz band
c	To cover bridge operation with 802.11 MACs (spanning tree)
d	To define physical layer requirements for 802.11 operation in other regulatory domains (countries)
e	To enhance 802.11 MAC for QoS
f	To develop recommended practices for Inter Access Point Protocol (IAPP) for multi-vendor use
g	To develop higher speed PHY extension to 802.11b (54 Mbps)
h	To enhance 802.11 MAC and 802.11a PHY-Dynamic Frequency selection (DFS), Transmit Power control (TPC)
i	To enhance 802.11 MAC security and authentication mechanisms
j	To enhance the 802.11 standard and amendments to add channel selection for 4.9 GHz and 5 GHz in Japan

Table 3-5 IEEE 802.11 Task Group Activities (continued)

k	To define RRM enhancements to provide interfaces to higher layers for radio and network measurements
k	To define Radio Resource Measurement enhancements to provide interfaces to higher layers for radio and network measurements
m	To perform editorial maintenance, corrections, improvements, clarifications, and interpretations relevant to documentation for 802.11 family specifications
n	Focus on high throughput extensions (>100MB/s at MAC SAP) in 2.4GHz and/or 5GHz bands
o	To provide Fast Handoffs in Voice over WLAN (goal is around 50ms)
p	Focus on vehicular communications protocol aimed at vehicles, such as toll collection, vehicle safety services, and commerce transactions via cars
r	To develop a standard specifying fast BSS transitions and fast roaming
s	To define a MAC and PHY for meshed networks that improves coverage with no single point of failure
t	To provide a set of performance metrics, measurement methodologies, and test conditions to enable manufacturers, test labs, service providers, and users to measure the performance of 802.11 WLAN devices and networks at the component and application level
u	To provide functionality and interface between an IEEE 802.11 access network (Hotspot) and any external network
v	To provide extensions to the 802.11 MAC/PHY to provide network management for stations (STAs)
w	To provide mechanisms that enable data integrity, data origin authenticity, replay protection, and data confidentiality for selected IEEE 802.11 management frames including but not limited to: action management frames, deauthentication and disassociation frames

RF Spectrum Implementations

In the United States, three bands are defined as unlicensed: industrial, scientific, and medical (ISM) bands. The ISM bands are as follows:

- 900 MHz (902-to-928 MHz)
- 2.4 GHz (2.4-to-2.4835 GHz) (IEEE 802.11b/g operates in this frequency range)
- 5 GHz (5.15-to-5.35 and 5.725-to-5.825 GHz) (IEEE 802.11a operates in this frequency range)

Each range has different characteristics. The lower frequencies exhibit better range, but with limited bandwidth and thus lower data rates. The higher frequencies exhibit less range and are subject to greater attenuation from solid objects.

The following sections cover some of the specific RF characteristics that the various 802.11 radios use for improving communications in the 2.4 and 5 GHz frequency ranges.

Direct Sequence Spread Spectrum

Direct sequence spread spectrum (DSSS) encodes redundant information into the RF signal. This provides the 802.11 radio with a greater chance of understanding the reception of a packet, given background noise or interference on the channel. Every data bit is expanded into a string of bits, or chips, called a chipping sequence or barker sequence. The chipping rate mandated by IEEE 802.11 is 11 chips per bit. It uses binary phase-shift keying (BPSK)/quadrature phase-shift keying (QPSK) at the 1 and 2 Mbps rates and 8 chips (complimentary code keying–CCK) at the 11 and 5.5 Mbps rate. This means that at 11 Mbps, 8 bits are transmitted for every one bit of data. The chipping sequence is transmitted in parallel across the spread spectrum frequency range.

IEEE 802.11b Direct Sequence Channels

14 channels are defined in the IEEE 802.11b direct sequence (DS) channel set. Each DS channel transmitted is 22 MHz wide, but the channel separation is only 5 MHz. This leads to channel overlap such that signals from neighboring channels can interfere with each other. In a 14-channel DS system (11 usable channels in the US), only three nonoverlapping (and thus, non-interfering) channels 25 MHz apart are possible (channels 1, 6, and 11).

This channel spacing governs the use and allocation of channels in a multi-AP environment, such as an office or campus. APs are usually deployed in a cellular fashion within an enterprise, where adjacent APs are allocated nonoverlapping channels. Alternatively, APs can be co-located using channels 1, 6, and 11 to deliver 33 Mbps bandwidth to a single area (but only 11 Mbps to a single client). The channel allocation scheme is illustrated in Figure 3-2.

Figure 3-2 IEEE 802.11 DSS Channel Allocations

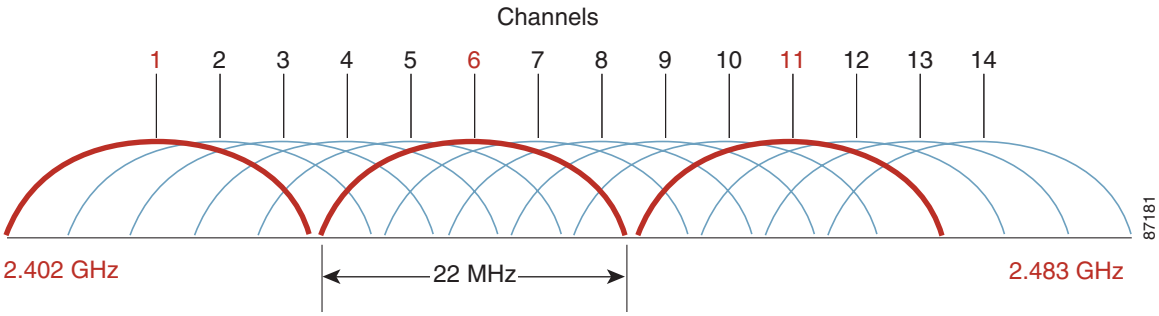

IEEE 802.11g

802.11g provides for a higher data rate (up to 54 Mbps) in the 2.4-GHz band, the same spectrum as 802.11b. 802.11g is backward-compatible with 802.11b and provides additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. At higher data rates, 802.11g uses the same modulation technique, orthogonal frequency division multiplexing (OFDM), as 802.11a (see IEEE 802.11a OFDM Physical Layer, page 3-9).

Table 3-6 lists 802.11g modulation and transmission types for the various data rates.

Table 3-6 802.11g Modulation and Transmission Types

Modulation	Transmission Type	Bits per Subchannel	Data Rate (Mbps)
BPSK	DSSS	NA	1
QPSK	DSSS	NA	2
CCK	DSSS	NA	5.5
BPSK	OFDM	125	6
BPSK	OFDM	187.5	9
CCK	DSSS	NA	11
QPSK	OFDM	250	12
QPSK	OFDM	375	18
16-QAM	OFDM	500	24
16-QAM	OFDM	750	36
64-QAM	OFDM	1000	48
64-QAM	OFDM	1125	54

IEEE 802.11a OFDM Physical Layer

IEEE 802.11a defines requirements for the physical layer of the OSI model, operating in the 5.0 GHz UNII frequency, with data rates ranging from 6 Mbps to 54 Mbps. It uses Orthogonal Frequency Division Multiplexing (OFDM), which is a multi-carrier system (compared to single carrier systems). OFDM allows subchannels to overlap, providing a high spectral efficiency. The modulation technique allowed in OFDM is more efficient than spread spectrum techniques used with 802.11b.

IEEE 802.11a Channels

The 802.11a channel shows the center frequency of the channels. The frequency of the channel is 10 MHz on either side of the dotted line. There is 5 MHz of separation between channels, as shown in Figure 3-3.

Figure 3-3 Channel Set

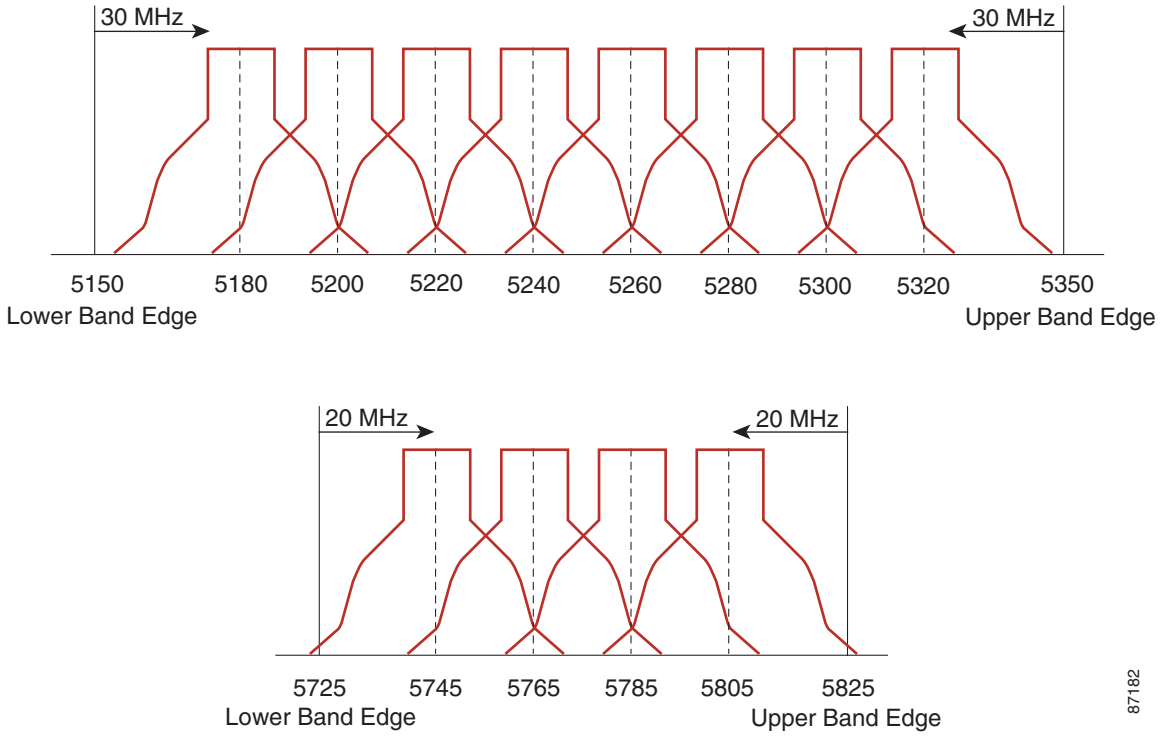

For the US-based 802.11a standard, the 5 GHz unlicensed band covers 300 MHz of spectrum and supports 12 channels. As a result, the 5 GHz band is actually a conglomerate of three bands in the USA: 5.150-to-5.250 GHz (UNII 1), 5.250-to-5.350 GHz (UNII 2), and 5.725-to-5.875 GHz (UNII 3).

RF Power Terminology

RF-specific terms, such as dB, dBi, and dBm are used to describe the amount of change in power transmitted measured at any discrete point, as perceived by the radio. The following sections cover their differences and provide a rule of thumb for their use, in addition to providing an explanation of effective isotropic radiated power (EIRP).

dB

The term *decibel* (*dB)* is mainly used for attenuation or amplification of the power level. dB is a logarithmic ratio of a signal to another standardized value. For example, dBm is where the value is being compared to 1 milliWatt, and dBw is where the value is being compared to 1 Watt.

The math is as follows:

Power (in dB) = 10 * log10 (signal/reference)

Plugging in some numbers (signal 100mW, reference 1mW) gives a value in dB of 20 (100 = 10 squared; taking the exponent 2 and multiplying by 10 gives you 20).

Remember that it is logarithmic (meaning that it increases or decreases exponentially and not linearly), and it is a ratio of some value to a reference. Also, remember that it is multiplied by 10.

Given that it is logarithmic, there are some general rules of thumb. An increase or decrease of 3 dB means that the signal doubled (double the power) or reduced in strength by 1/2, respectively. An increase or decrease of 10dB means that the signal went up by 10 times or down to $1/10^{th}$ the original value.

Indoor WLAN and outdoor WLAN deployments both offer separate challenges in RF deployments, and need to be analyzed separately. However, there are some rules of thumb for indoor use. For every increase of 9dB, the indoor coverage area should double. For every decrease of 9dB, the indoor coverage area should be cut in half.

dBi

The term *dBi* is used to describe the power gain rating of antennas. The real antennas are compared to an isotropic antenna (a theoretical or imaginary antenna) that sends the same power density in all directions, thus the use of dBi.

Antennas are compared to this ideal measurement, and all FCC calculations use this measurement (dBi). For example, a Cisco omni-directional AIR-ANT4941 has a gain of 2.2 dBi, meaning that the maximum energy density of the antenna is 2.2 dB greater than an isotropic antenna.

dBm

The term *dBm* uses the same calculation as described in the dB section, but has a reference value of 1 milliWatt.

So, taking into consideration the example previously given in the dB section, if the power jumped from 1 mW to 100mW at the radio, the power level would jump from 0 dBm to 20 dBm.

Besides describing transmitter power, dBm can also describe receiver sensitivity. Receiver sensitivity is in minus dBm (-dBm), because the signal reduces in value from its point of transmission. The sensitivity indicates the lowest power the receiver can receive before it considers the signal unintelligible. The approximate receiver sensitivity of Cisco radios is -84dBm for a 1200 series “a” radio and -90 dBm for the 1200 “g” radio.

Note The “g” radio can recover a signal at half the strength of the “a” radio. It is generally true that for a given technology, lower frequency radios can achieve a better sensitivity than higher frequency radios.

Effective Isotropic Radiated Power

Although transmitted power based on the radio setting is rated in either dBm or Watts, the maximum energy density coming from an antenna from a complete system is measured as effective isotropic radiated power (EIRP), which is a summation of the dB values of the various components. EIRP is the value that regulatory agencies, such as the FCC or ETSI, use to determine and measure power limits, expressed in terms of maximum energy density within the first Fresnel of the radiating antenna. EIRP is calculated by adding the transmitter power (in dBm) to antenna gain (in dBi) and subtracting any cable losses (in dB). For example, if you have a Cisco Aironet 350 bridge connected to a solid dish antenna by a 50 foot length of coaxial cable, plugging in the numbers gives the following:

- Bridge—20 dBm
- 50 Foot Cable—3.3 dBm (negative because of cable loss)
- Dish Antenna—21 dBi
- EIRP—37.7dBm

For more information, see the following URL:

http://www.cisco.com/en/US/partner/tech/tk722/tk809/technologies_tech_note09186a00800e90fe.shtml

Planning for RF Deployment

Many of the RF-design considerations are interdependent or implementation-dependent. As a result, there is no “one-size-fits-all” template for the majority of requirements and environments.

Different Deployment Types of Overlapping WLAN Coverage

How much overlapping WLAN coverage you set in your wireless network depends on the usage, though with limited exceptions, all designs should be deployed to enable low latency coverage. Wireless networks can be deployed for location management, voice, or data-only networks, or a combination of all three. The difference is in the pattern in which the APs are laid out, and the amount of RF overlap in the coverage area. When planning a WLAN deployment consideration should be given to future uses of the WLAN deployment.

Converting a WLAN deployment to support additional services beyond a data-only deployment is not simply a matter of adding APs; it can require an additional site survey and the possible relocation of existing APs.

Data-Only Deployment

Data-only deployments do not require a large amount of overlap. This is because 802.11 clients respond to a lower signal from a nearby AP, should one fail, by stepping down their rate and taking a longer time to transmit. The required overlap is determined by the WLAN data rate requirement described in WLAN Data Rate Requirements, page 3-15. Minimal overlap is required for data-only networks, which allow all data rates. For data-only networks, the rule of thumb for separation of APs is typically 120–130 feet, but, when making your estimation for AP separation, remember to factor in objects that affect RF coverage, such as wall densities, machinery, elevators, or even wide-open space with steel cages, because your results can vary depending on the RF environment. RRM has been developed for this type of deployment and it is very useful for controlling the RF coverage.

Voice/Deployment

Figure 3-4 shows the voice network pattern and overlap. The APs are grouped closer together and have more overlap than a data-only installation, because voice clients should roam to a better AP before dropping packets. You generally also want to run smaller cells than in the past, and ensure the overlapping cell edges at or above -67 dBm. This accomplishes a number of things including greater homogeneity across a single cell and reducing processor load in the handheld, which increases link stability and reduces latency. Although only one AP might be required for a defined area, Cisco recommends that you have two APs on nonoverlapping channels with a received signal strength indication (RSSI) above 35 at all times in your installation, for redundancy and load balancing purposes. For the 7920 voice deployment, Cisco recommends that you have a Received Signal Strength Indication (RSSI) above 35 at all time in your installation. This is to ensure that the VoIP phone has good reception as well as allowing some over-subscription and enhances roaming choices for the phone.

Remember that designing for low noise background is as important as relatively high energy density within the cell. This means that a good baseline power setting for the AP is in the 35–50 mW range. This generally requires approximately 15 percent more APs than if you deployed a coverage model at 100 mW. The smaller the cell, generally the more homogenous the energy density across the cell.

100 mW designs also generate significantly greater noise load into a given area; 35–50 mW designs are commonly 10 dBm or lower in background noise, even in complex coverage areas.

Pre-site surveys are useful for identifying and characterizing certain challenging areas and potential sources for interference, such as existing WLANs, rogues, and non-802.11 interference from sources such as microwave ovens and cordless telephones. Following a design that should be reviewed and approved by all stakeholders, post-site surveys should be considered as an excellent audit mechanism to ensure that the coverage model complies with the intended functional requirements as set forth by the stakeholders.

When making your estimation for separation, remember to factor in objects that affect RF coverage such as wall densities, machinery, elevators, or even wide open spaces with steel cages, because your results may vary depending on the RF environment. Be sure to include transient dynamics such as forklifts, large groups of people, or large objects moved through the area by crane or similar load bearing devices. Characterizing for other wireless deployments such as robotic systems, cranes, and so on, are also a good idea. A WLC is often a very effective method for preliminary site evaluation, by allowing a fast deployment of a WLAN infrastructure that can then be used to make RF measurements of the area; a hand-walked site survey is also effective insurance for complex areas such as those commonly found in healthcare, retail, and manufacturing.

Figure 3-4 Single Floor Site Survey for Voice

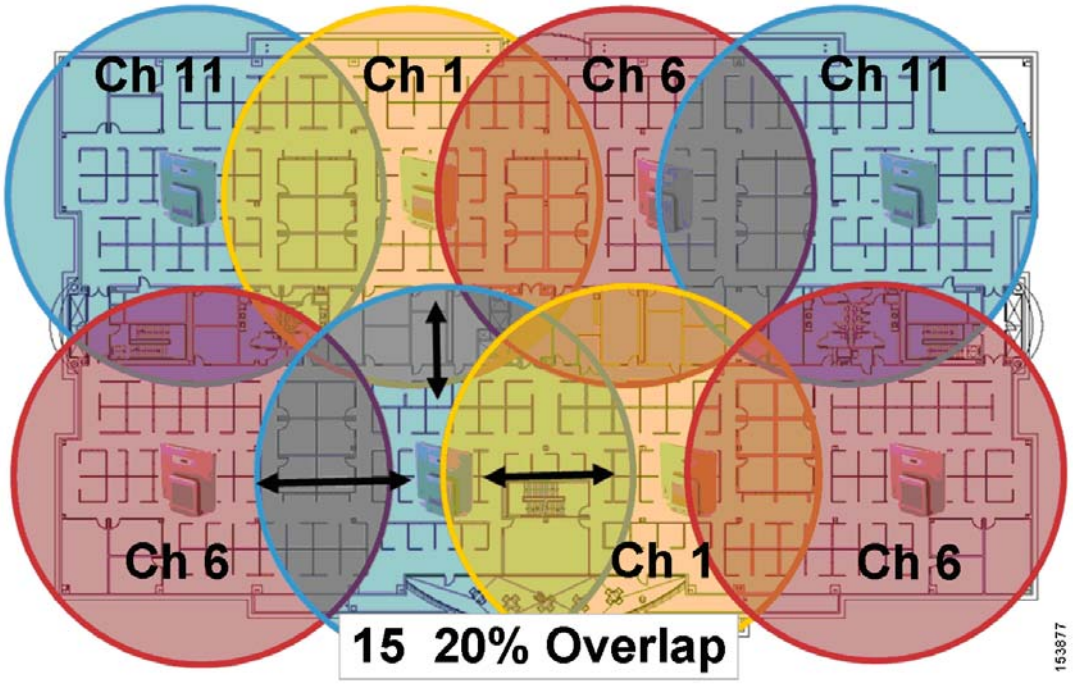

For more information on a wireless voice deployment, see Chapter 11, "VoWLAN Design Recommendations," as well as the 7920 deployment guide at the following URL: http://www.cisco.com/en/US/partner/products/hw/phones/ps379/products_implementation_design_guide_book09186a00802a029a.html

For helping configuring your lightweight APs and controller for 7920 voice operations, see the following URL:
http://www.cisco.com/en/US/partner/products/ps6366/prod_technical_reference09186a00805e75a1.html

Location-Based Services Deployments

The third type of deployment is the location-based services (LBS) deployments, which may be the most complex of current applications because it relies not only on excellent cell coverage, but optimal location of APs. Location management deployments can simultaneously track thousands of devices by using the WLAN infrastructure. Examples include Wi-Fi tag type deployments or asset tracking deployments to locate equipment or devices via the wireless network and/or simply to indicate where wireless clients are throughout the wireless network in relation to a drawing or diagram. This can be used to make the wireless infrastructure more secure by providing the location of a rogue client or APs, and greatly improve client troubleshooting capabilities.

For a location management deployment, the APs are laid out in a staggered pattern. Figure 3-5 shows a typical pattern. The staggered pattern allows for more accurate estimation of the location of a device.

For a discussion of Location-Based Services, see Chapter 15, "Cisco Unified Wireless Location-Based Services," and the white paper entitled *WiFi Location Based Services—Design and Deployment Considerations*, which can be found at http://www.cisco.com.

Figure 3-5 Example of a Single Floor Location Management Deployment

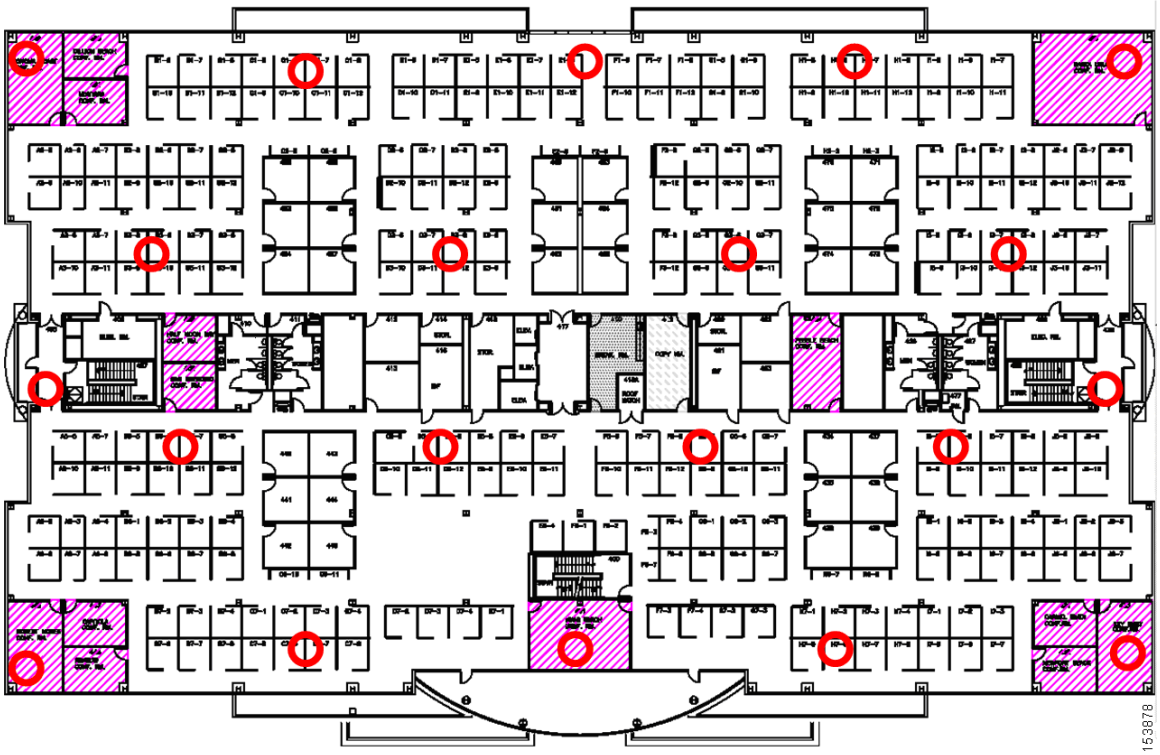

WLAN Data Rate Requirements

Data rates affect AP coverage areas. Lower data rates (such as 1 Mbps) can extend the coverage area farther from the AP than higher data rates (such as 54 Mbps) as illustrated in Figure 3-6 (which is not drawn to scale). Therefore, the data rate (and power level) affects coverage and consequently the number of APs required for the installation, as illustrated in Figure 3-7 for different data rates. As part of the planning process, consider the required data rates, the required range, and the required reliability.

Data Rate Compared to Coverage Area

Different data rates are achieved by the AP sending a redundant signal on the wireless link, allowing data to be more easily recovered from noise. The number of symbols, or chips, sent out for a packet at the 1 Mbps data rate is greater than the number of symbols used for the same packet at 11 Mbps. This means that sending data at the lower bit rates takes more time than sending the equivalent data at a higher bit rate. And when there is more than one client associated to the radio, the lower rate client affects the higher rate clients' maximum data rate by taking longer to transmit a packet of the same length.

Figure 3-6 Data Rate Compared with Coverage

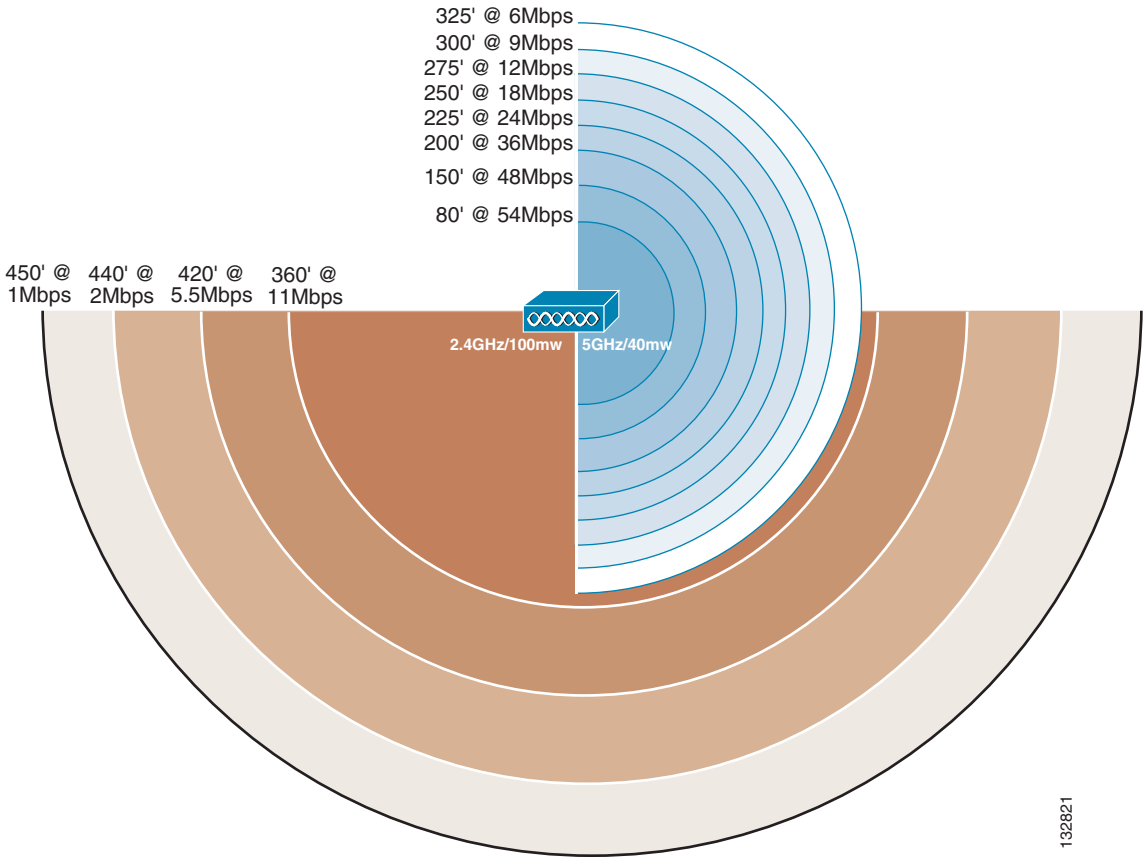

The actual diameter of the coverage depends on factors such as environment, power level, and antenna gain. For example, indoors using the standard antennas on the NIC card and APs, the diameter of the 1 Mbps circle is approximately 700 feet (210 m), and the diameter of the 11 Mbps circle is about 200 feet (60 m). Increasing the gain of the antenna can increase the distance and change the shape of the radiation pattern to something more directional.

AP Density for Different Data Rates

The minimum required reliable data rate has a direct impact upon the number of APs needed in the design, along with power setting, antenna gain, and location. Figure 3-7 shows coverage comparison and AP density for different data rates. Although six APs with a minimum data rate of 2 Mbps might adequately service an area, it might take twice as many APs to support a minimum data rate of 5 Mbps, and more again to support a minimum data rate of 11 Mbps for the same coverage area.

Figure 3-7 Coverage Comparison and AP Density for Different Data Rates

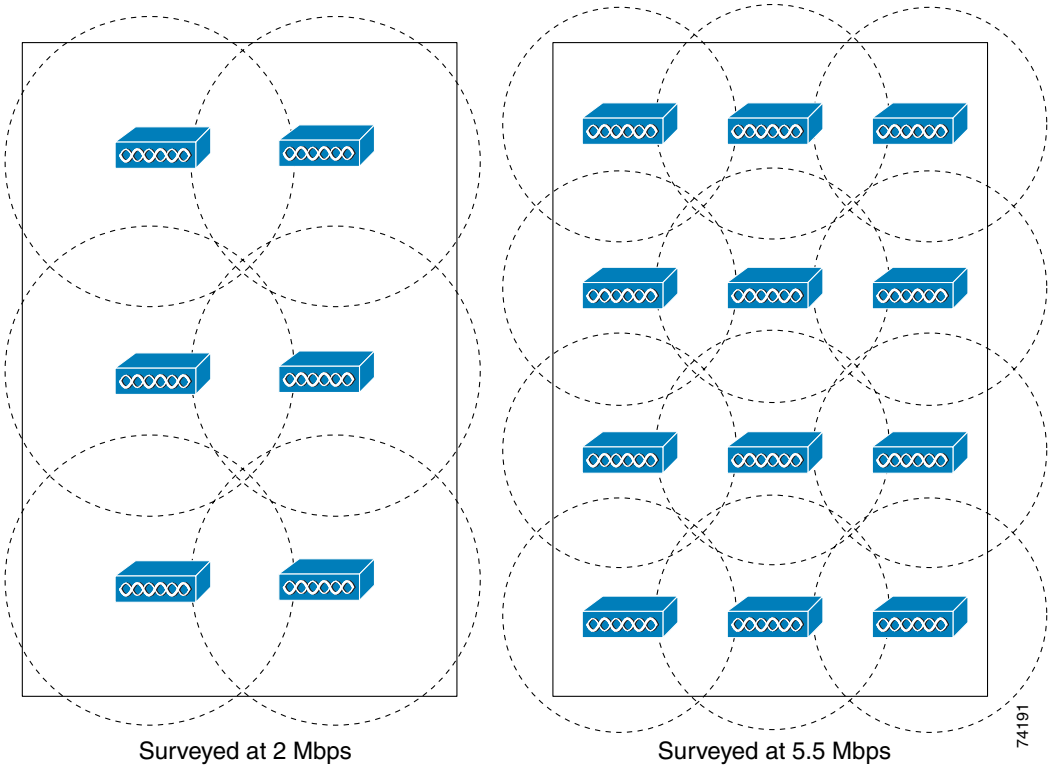

The data rate you choose depends on the type of application to be supported, but should not be greater than the typical requirements because there is tradeoff in coverage. In a typical WLAN environment, the higher data rates give maximum throughput and should minimize performance-related support issues. In a WLAN vertical application environment, the data rates selected are determined by the application requirements of speed and reliability as measured by delay and jitter. The physical facility and/or whether the network is client-centric generally dictates range requirements; some clients might not support the higher data rates, longer ranges, or the delay and jitter rates of an infrastructure element such as an AP.

It might seem logical to choose the default configuration of APs and clients, thereby allowing all data rates. However, there are three key reasons for limiting the data rate to the *highest* rate at which full coverage is obtained:

- Broadcast and multicast (if enabled) are sent at the lowest associated data rate (to ensure that all clients can receive the packets). This reduces the throughput of the WLAN because traffic must wait until frames are processed at the slower rate.

- Clients that are farther away, and therefore accessing the network at a lower data rate, decrease the overall throughput by causing delays while the lower bit rates are being serviced. It might be better to force the clients to roam to a closer AP so as not to impact the performance of the rest of the network.
- If a 54 Mbps service is specified and provisioned with APs to support *all* data rates, clients at lower rates can associate with the APs that can create a coverage area greater than planned, thereby increasing the security exposure (by allowing association from outside the building) and potentially interfering with other WLANs.

Client Density and Throughput Requirements

Wireless APs have two characteristics that make actual client data throughput slower than the data rate:

- APs have an aggregate throughput less than the data rate because 802.11 provides a reliable transport mechanism that ACKs all packets, thereby halving the throughput on the channel.
- APs are similar to shared hubs. That is, the channel is shared by all the clients associated to that AP on that channel, thus collisions slow data throughput.

With this in mind, you must have the approximate estimate of the maximum number of active associations (active clients). This can be adjusted more or less according to the particular application.

Each cell provides an aggregate amount of throughput that is shared by all the client devices that are within the cell and associated to a given AP. This basically defines a cell as a collision domain. After deciding on the minimum data rate, be sure to consider how much throughput should, on average, be provided to each user of the wireless LAN.

Take the example of barcode scanners; 25 Kbps may be more than sufficient bandwidth for such an application because using an 802.11b AP at 11 Mbps of data rate results in an aggregate throughput of 5–6 Mbps. A simple division results in a maximum number of 200 users that can theoretically be supported. This number cannot in fact be achieved because of the 802.11 management overhead associated with the large number of clients and packet collisions. For a 1 Mbps system, 20 users can use the same AP for similar bandwidth results.

You can increase the potential per-user throughput by decreasing the number of users contending for the aggregate throughput provided by a single AP. This can be done by decreasing the size of the coverage area, or adding a second AP on a non-overlapping channel in the same coverage area. To reduce the coverage area, the AP power or antenna gain can be reduced, resulting in fewer clients in that coverage area. This means you need more APs for the same overall area, increasing the cost of deployment. An example of this is shown in Figure 3-8.

Figure 3-8 Changing the Output Power to Increase Client Performance

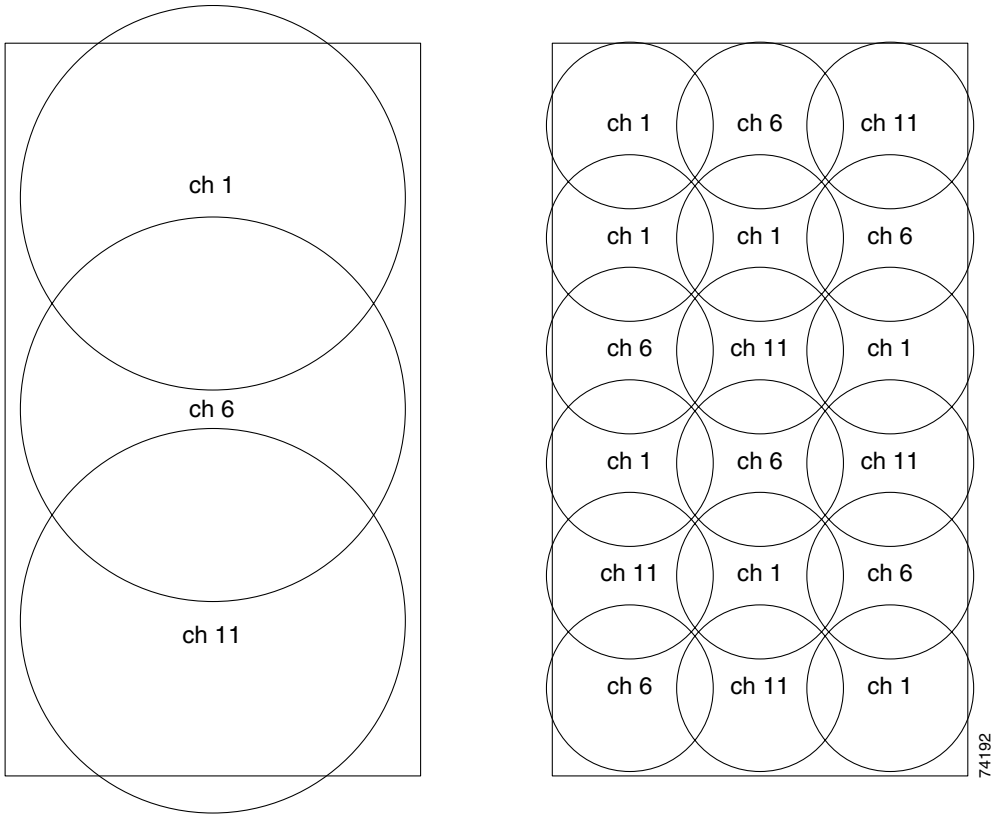

180 Users per floor
30 mW transmitter power
3 Accss Points
60 users per AP
11 Mbps data rate

180 Users per floor
5 mW transmitter power
18 Accss Points
10 users per AP
11 Mbps data rate

Note Client power should be adjusted to match the AP power settings. Maintaining a higher setting on the client does not result in higher performance and it can cause interference in nearby cells.

WLAN Coverage Requirements

Different enterprises have different coverage requirements. Some need a WLAN to cover specific common areas, while others need WLANs to cover each floor of a building or to cover the entire building including stairwells and elevators, or to cover the entire campus including car parks and roads. Apart from impacting the number of APs required, the coverage requirements can introduce other issues, such as specialized antennas, outdoor enclosures, and lightning protection.

Power Level and Antenna Choice

Power level and antenna choice go hand-in-hand to determine AP placement. Together, these two variables determine where and how powerful the RF is in any given place in the environment. Along with choosing the correct antenna to produce the required coverage area, Cisco recommends the use of RRM to control the power level and provide the optimal channel/power plan through the use of special reuse (see Radio Resource Management (Auto-RF), page 3-28 for more information).

To understand antenna principles a little better, a brief introduction to the three fundamental properties follows:

- An antenna gives the wireless system three fundamental properties: gain, direction, and polarization. Gain and direction mandate range, speed, and reliability; polarization affects reliability and isolation of noise.

  Gain is a measure of increase in power. Direction is the shape of the transmission pattern. A good analogy for an antenna is the reflector in a flashlight. The reflector concentrates and intensifies the light beam in a particular direction similar to what a parabolic dish antenna does to an RF source in a radio system. RF has an electric field component and a magnetic field component that are typically confined to a plane perpendicular to the propagation direction. Polarization indicates the direction of the electric field; if the electric field is orientated vertically, the wave is said to be vertically polarized. Polarization can be simply related to how the antenna is positioned. An antenna that is parallel to the Earth's surface produces horizontally polarized radio waves.

Omni and Directional Antennas

Omni antennas have a different radiation pattern compared to isotropic antennas that are theoretical. The omni antenna features a radiation pattern that is nearly symmetric about a 360 degree axis in the horizontal plane, and 75 degrees in the vertical plane (assuming the dipole antenna is standing vertically). The radiation pattern of an omni antenna generally resembles a donut in shape.

Regarding antenna choice, you must consider the RF pattern produced by the antenna because the type of antenna (omni or directional) affects RF coverage by focusing the bulk of the RF energy in a specific direction, pattern, and density.

For example, with an omni-directional antenna, the antenna radiates in a donut pattern. Figure 3-9 shows an omni-directional antenna RF radiation pattern in the vertical and horizontal direction. This is an actual measurement, so it does not follow the donut lines perfectly, but does show from where the donut pattern comes. As described above, other RF-affecting variables (people in the room, amount of devices stored in the facility, leaves on trees for outdoor deployment, interference from different RF sources, and so on) may affect the real RF coverage pattern.

Figure 3-9 Omni-Directional RF Pattern

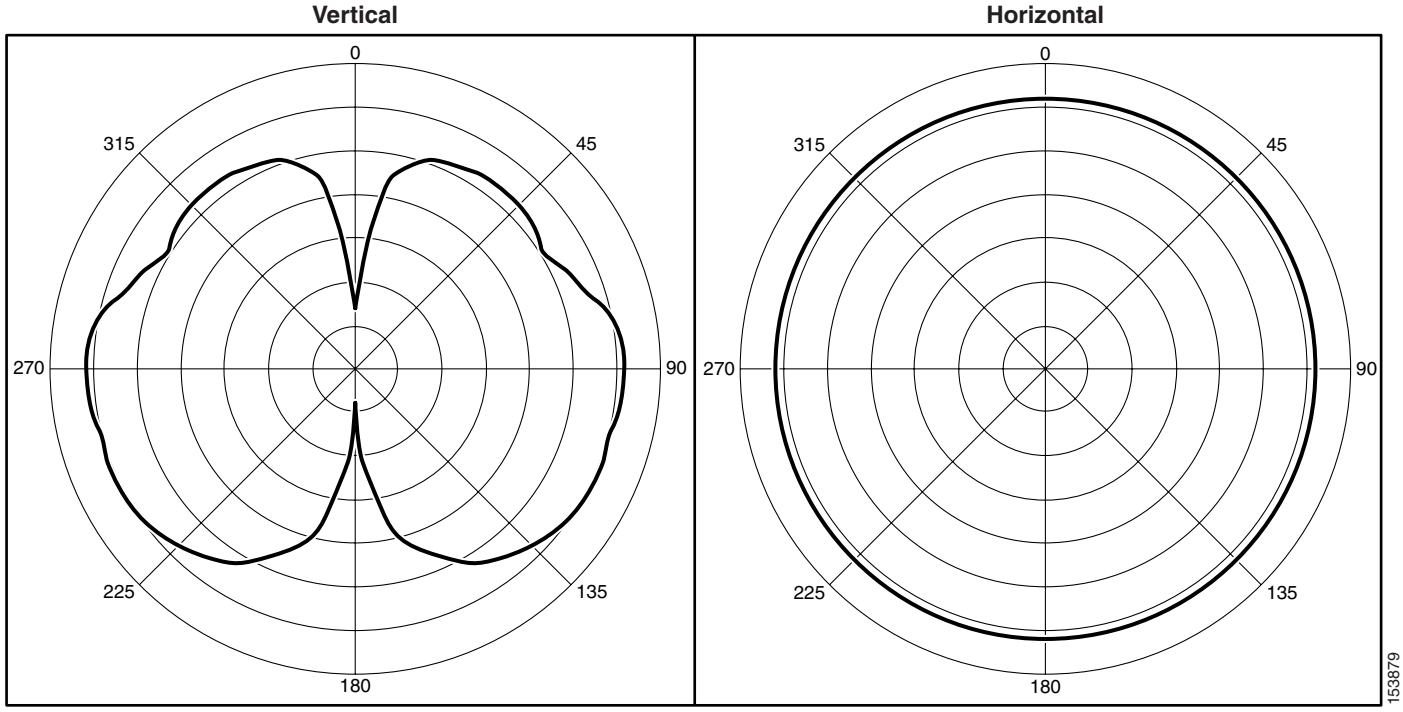

Looking at the pattern in Figure 3-9, this may be the incorrect antenna to use on a wall, especially if it is mounted along an exterior wall where the pattern can radiate outside of the building. This can open up the wireless network to hackers outside the building and compromise the wireless network.

Patch Antennas

Another type of directional antenna is a patch antenna. Patch antennas not only radiate away from the wall or place where they are mounted, but also have rear and side lobes that produce a weakened but still quite effective RF region. Figure 3-10 shows the real horizontal pattern of a diversity patch wall mount antenna. Although most of the coverage area is in front of the patch antenna, notice the back and side RF pattern from the center area. Again, antenna selection is important because it defines the radiation pattern and where wireless connectivity is possible.

Figure 3-10 Patch Wall Mount Antenna Horizontal Plane

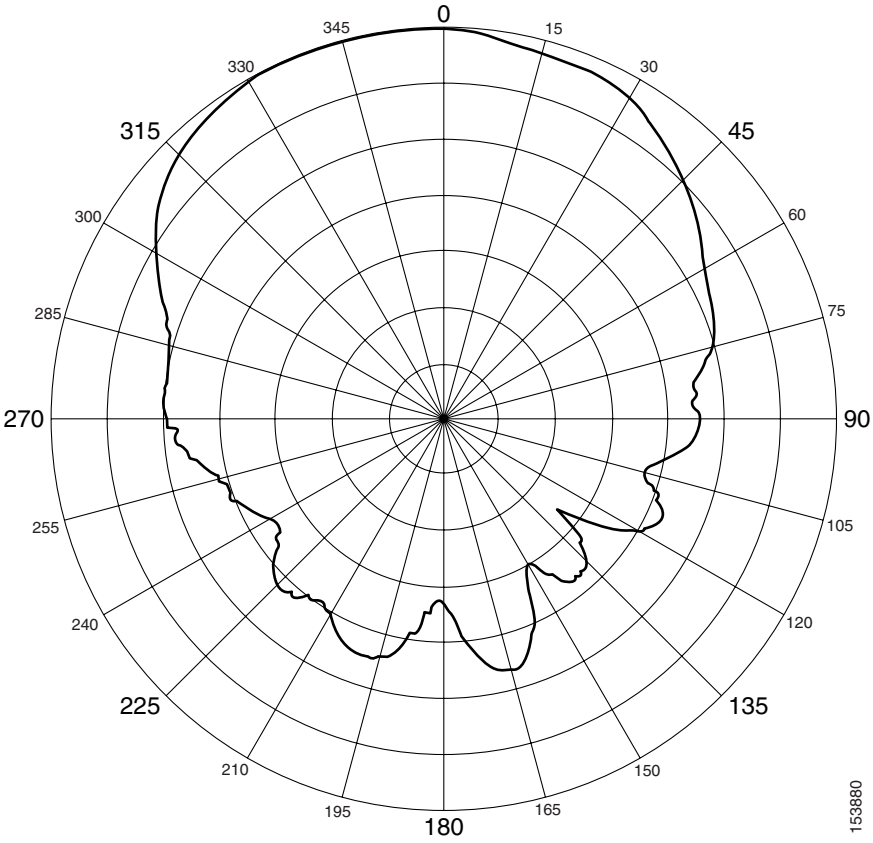

For more information on antenna selection, see the *Cisco Antenna Selection Guide* at the following URL:

http://www.cisco.com/en/US/products/hw/wireless/ps430/products_data_sheet09186a008008883b.html

Security Policy Requirements

A good RF design can effectively minimize unintended RF radiation in areas not requiring coverage. For example, if WLAN coverage is required only in buildings and not outside, then the amount of RF coverage outside of the buildings can be minimized by using the correct power setting, AP placement and directional antennas pointing inwards towards the center of the building or areas. By tuning RF transmit levels and using the correct antenna for the coverage area, you can reduce the amount of RF that radiates outside the buildings to decrease the security exposure. This can reduce the exposure of wireless network to hackers outside the building or coverage area, and avoid a compromise of the wireless network.

RF Environment

The performance of the WLAN and its equipment depends on its RF environment, equipment, selection, coverage design, quality of audits, configurations, and quality of deployment. The following are some examples of adverse environmental variables that can disrupt wireless communications by either providing interference on the channel or in some way changing the RF characteristics of the signal:

- 2.4 GHz cordless phones
- Walls fabricated from wire mesh and stucco
- Filing cabinets and metal equipment racks
- Transformers
- Heavy duty electric motors
- Fire walls and fire doors
- Concrete
- Refrigerators
- Sulphur plasma lighting (Fusion 2.4 GHz lighting systems)
- Air conditioning duct-work
- Other radio equipment
- Microwave ovens
- HVAC ducting
- Large transient elements such as forklifts or metal fabrications
- Other WLAN equipment

A site survey might be required to ensure that the required data rates are supported in all of the required areas, often driven by the environmental variables mentioned above, although a WLC is an excellent resource for site pre-planning and initial identification of RF challenges as well as channel and power settings.

The site survey should also consider the three dimensional space occupied by the WLAN. For example, a multi-story building WLAN with different subnets per floor might require a different RF configuration than the same building with a single WLAN subnet per building. In the multiple subnet instances, a client attempting to roam to a different AP on the same floor might acquire an AP from an adjacent floor, which is not a desired behavior. Switching APs in a multi-subnet environment changes the roaming activity from a *seamless Layer 2 roam* to a *Layer 3 roam*, which in turn might disrupt sessions and require user intervention. This Layer 3 roaming in WLCs simplifies this deployment by allowing the client to maintain its IP address even if the AP to which it roams is on another subnet.

RF Deployment Best Practices

Some design considerations can be addressed by general best practice guidelines. The following applies to most situations:

- The number of users versus throughput and a given AP. A common recommended number of users per AP is 15 to 25 for data-only users only and, for the 7920 VoIP (or similar voice devices) wireless handset, 7 to 8 voice users when data is present. This number should be used as a guideline and may vary depending on the handset in use. Check your handset requirements.
- The distance between APs can cause throughput variations for clients, based on the distance from the AP. The recommendation is to limit the AP data rate to the higher data rates to minimize data rate shifting.
- The number of APs depends on coverage and throughput requirements, which can vary. For example, Cisco System's internal information systems (IS) group currently uses six APs per 38,000 square feet of floor space for data-only operation.

Note Based on the variability in environments, Cisco recommends that a site survey be performed to determine the number of APs required and their optimal placement.

Manually Fine-Tuning WLAN Coverage

A number of factors can affect the WLAN coverage, as follows:

- Channel and data rate selection
- Overlapping WLAN coverage for location management, voice, or data-only
- Power level
- Antenna choice (omni-directional, or patch)

For a given data rate and location, the WLAN designer may alter power levels and/or elect to use a different antenna, to effect changes to the coverage area and/or coverage shape. Altering power levels or channel selection can be done manually as described below, or the Cisco Wireless Controller can do this automatically via the Radio Resource Management (RRM) algorithms, also referred to as Auto-RF. Cisco recommends the use of Radio Resource Management (RRM) to control the power level and channel, keeping in mind that the channel changing algorithm is highly dampened so that only a very disruptive (and persistent) interference source would cause a change to the channel topology, which in turn would cause clients to reassociate and any voice calls to be dropped. Changes in AP power do not impact clients. (See Radio Resource Management (Auto-RF), page 3-28 for more details).

Channel and Data Rate Selection

Channel selection depends on the frequencies that are permitted for a particular region. For example, the North American and ETSI 2.4 GHz channel sets permit allocation of three nonoverlapping channels: 1, 6, and 11 while the 5 GHz channel set permits 12 channels.

The channels should be allocated to the coverage cells as follows:

- Overlapping cells should use nonoverlapping channels.
- Where channels must be re-used in multiple cells, those cells should have minimal overlap with each other. Figure 3-11 shows this pattern.

Recommendations for Channel Selection

Channel selection can be done manually, as described below.

Figure 3-11 Channel Allocated To APs

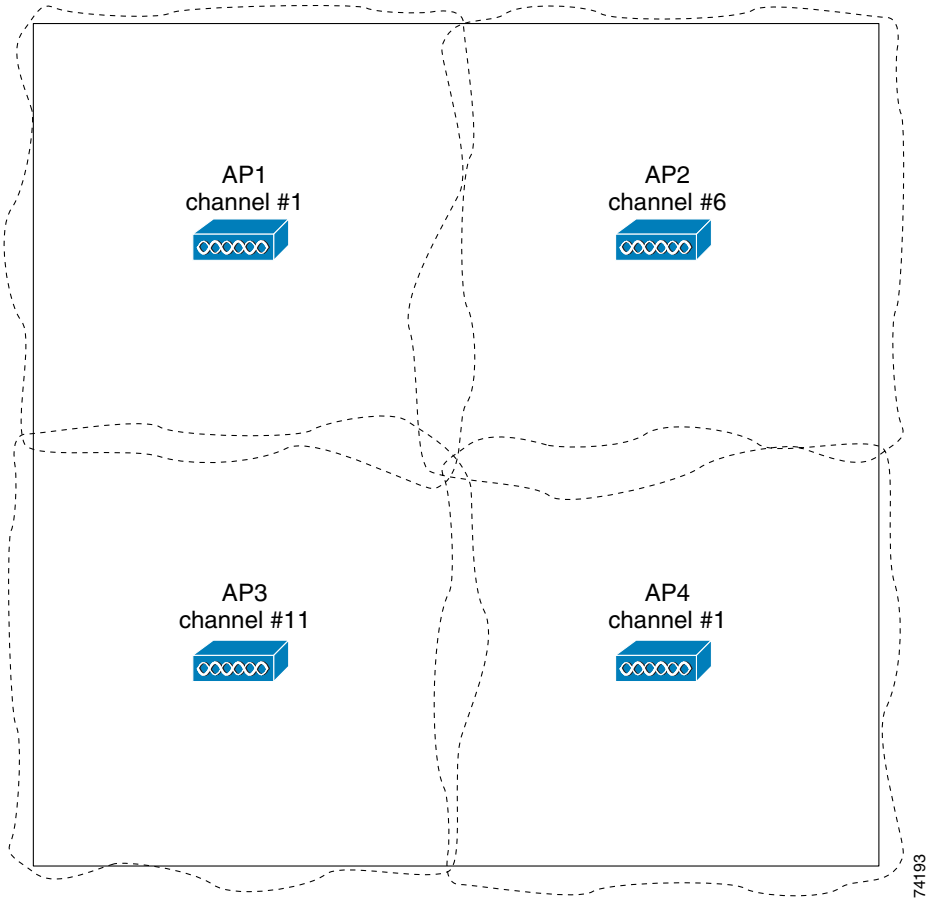

A site survey should be conducted using the same frequency plan as intended for the actual deployment. Some sites have high noise backgrounds which may prohibit the use of one or more channels. This provides a better estimate of how a particular channel at a particular location will react to the interference and the multipath. Channel selection also helps in planning for co-channel and the adjacent channel interference, and provides information about where you can reuse a frequency (see Figure 3-12).

In multi-story buildings, check the cell overlap between floors, especially where windows may be located, according to these rules/guidelines. Careful pre-planning and selection of AP location might be required in approximately 10 percent of the cases. Multi-story structures such as office towers, hospitals, and university classroom buildings introduce a third dimension to coverage planning. The 2.4 GHz waveform of 802.11b and 802.11g can pass through many walls. The 5 GHz waveform of 802.11a has approximately half the tendency for a given power to transmit suitable amounts of energy through walls because of its higher frequency. With 2.4 GHz Wi-Fi LANs in particular, you must not only avoid overlapping cells on the same floor, but also on adjacent floors when coverage models include cells that cover windows on both floors. With only three channels, this can be achieved through careful three-dimensional planning.

As a final step, after setting up the WLAN network, you should always retest the site using the selected channels and check for any interference. Keep in mind that the RRM algorithms are logical and subject to the physical topology of the network. It thus takes into account the three-dimensional placement of APs and provides the optimal channel/power setting for the sampling interval.

Manual Channel Selection

Figure 3-12 shows a screenshot of the web page for configuring one of the 802.11b/g radios under the wireless selection. On the top right-hand side, channel 11 has been manually selected and the transmit power is set to 1, the highest level (8 sets the AP to the lowest level).

Note The assignment method should normally be left at the global setting, unless there is a desire to manually control these settings. This allows the controller to dynamically change the channel number and transmit power as determined by the RRM. See Radio Resource Management (Auto-RF), page 3-28 for more information.

Figure 3-12 Channel Assignment

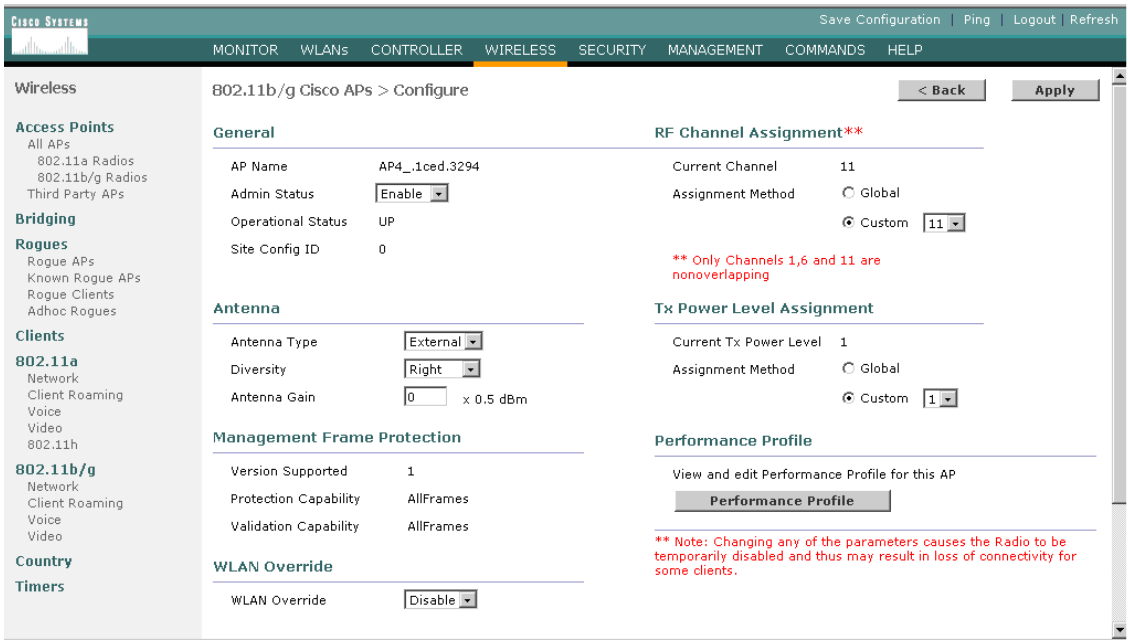

It is also possible to implement a dual-band deployment scheme, as shown in Figure 3-13. The top left portion of the diagram shows the 802.11b/g-only deployment, which uses the three nonoverlapping channels (channels 1, 6, and 11) to map out a pattern that has the least co-channel interference; that is, interference from an AP close by that is on the same channel, that is operating at sufficient power levels with its coverage pattern overlapping with that of another access point. It also shows an 802.11a deployment, which uses the eight nonoverlapping channels. The right side of the diagram illustrates how the channels would be mapped in a dual-band deployment.

Figure 3-13 Dual Band Deployment Diagram

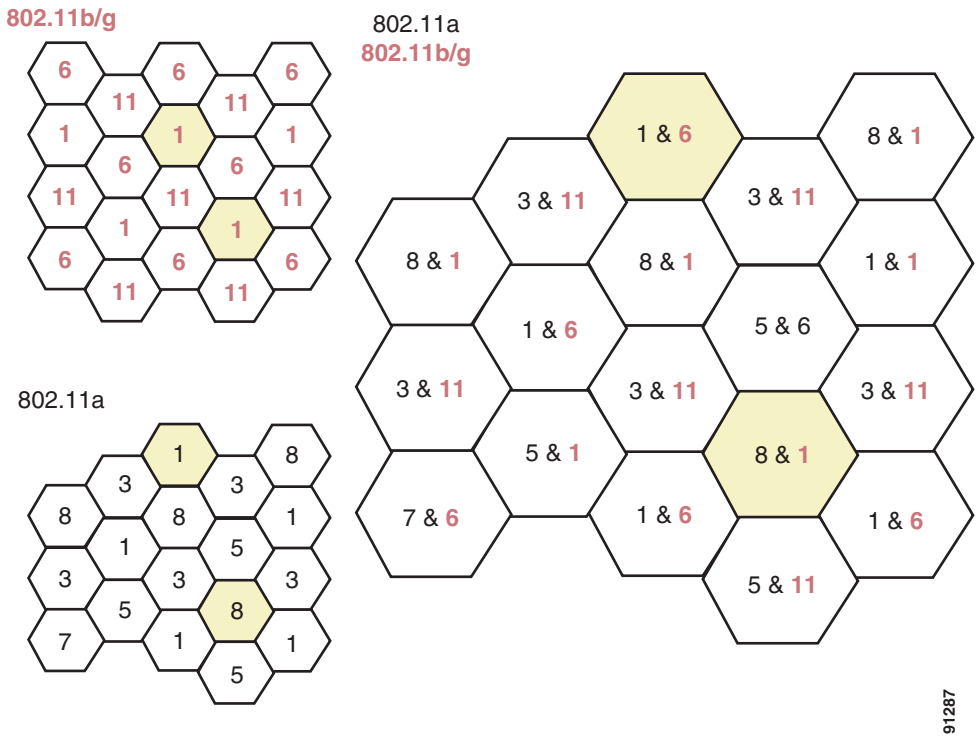

Data Rate Selection

Figure 3-14 is a screenshot of the web page of the global 802.11b/g parameters. The data rate settings are shown on the right side of the screen.

Figure 3-14 Data Rate Assignment

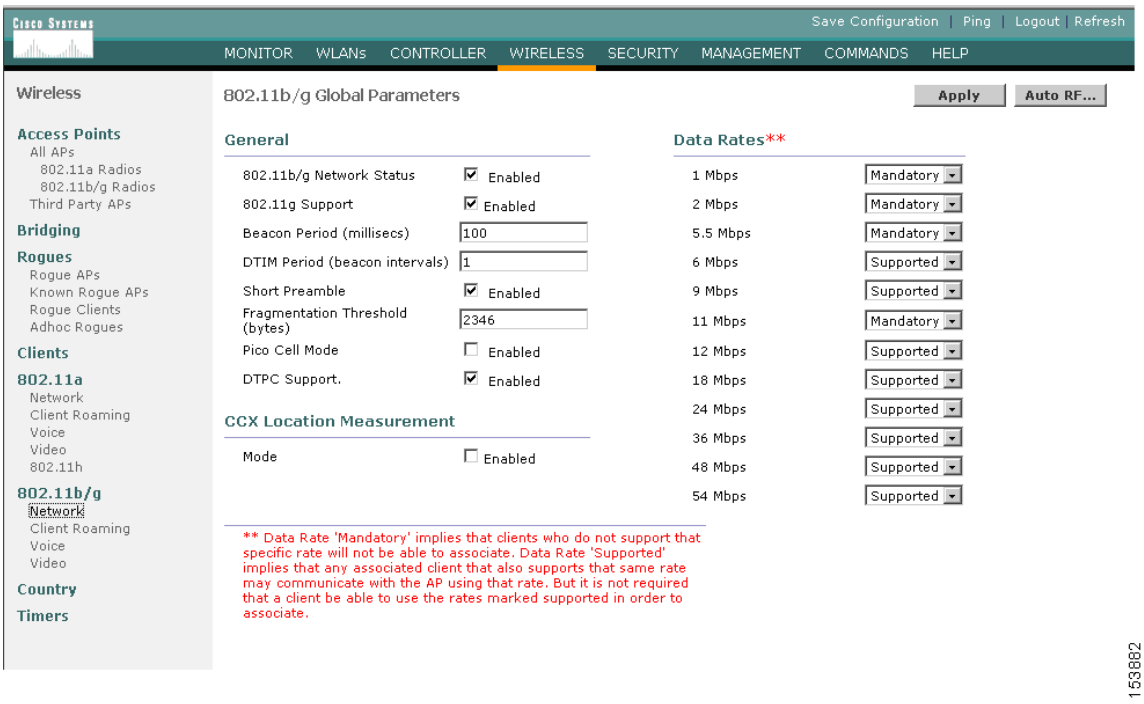

Mandatory, Supported, and Disabled Rate Modes

You can use the data rate settings to choose which data rates the wireless device can use for data transmission. There is a direct correlation between data rates, range, and reliability. The lower the data rate, the greater the reliability and range for a given power setting. Sites vary for specifics, but a reasonable rule of thumb for carpeted space is an order of magnitude of increased reliability for every time you halve the data rate. Range is generally affected by a factor of a 30 percent increase (approximately) for every halving of data rate. Managing the square footage of the area covered within a -67 dBm edge can be effectively managed using this technique. Setting the data rates to match client, application, or user needs is an effective RF design element that should be considered before deploying APs.

Data rates are expressed in megabits per second. You can set each data rate to one of three modes:

- Mandatory—Allows transmission at this rate for all packets, both unicast and multicast. The data rate on at least one of the APs must be set to Mandatory, and all clients that associate to the AP must be able to physically support this data rate on their radio to use the network. Additionally, for the wireless clients to associate to the AP, they must be able to currently receive packets at the lowest mandatory rate and their radios must physically support the highest mandatory data rate. If more than one data rate is set to mandatory, multicast and broadcast frames are sent at the highest common mandatory transmission rate of all associated clients (the lowest mandatory receive rate of all of the clients). This allows all clients to receive broadcast packets. The lowest mandatory rate is normally set at 1 Mb/s.
- Supported—Allows transmission at this rate for unicast packets only. The AP transmits only unicast packets at this rate; multicast and broadcast packets are transmitted at one of the data rates set to mandatory. The wireless clients always attempt to transmit and receive at the highest possible data rate. They negotiate with the AP for the highest data rate set to supported or mandatory to transmit and receive unicast packets. The wireless client devices are able to receive broadcast or multicast packets at any mandatory rate at or below the negotiated rate.

- Disabled—The AP does not transmit data at this rate.

Lowest and Highest Mandatory Rate Settings

Multiple clients associated to the AP can have completely different transmission rates, depending on interference, obstacles, or their distance from the AP. For example, if an 802.11b client is far from the AP and can only transmit and receive at a speed of 1 Mb/s because of the distance, it would be able to associate to the AP because the lowest mandatory rate (see Figure 3-14) is set to 1 Mb/s. If a second 802.11g client associates to the AP at 54 Mb/s, the AP would transmit broadcasts and multicasts at 1 Mb/s because this is the highest mandatory rate that all clients can receive. If the lowest mandatory rate was set to 5.5 Mb/s, the 802.11b client would not be able to associate to the AP because it could not receive broadcast packets at the lowest mandatory rate.

In Figure 3-14, note that the highest mandatory setting is 11 Mb/s. The highest mandatory rate tells the AP what rate the client radios must be able to physically transmit at. This does not mean that they are actually transmitting and receiving packets at that rate, it just means that the radio physically supports that rate; the wireless client needs only to be able to receive packets at the lowest mandatory rate. 802.11b devices would be able to associate to the AP shown in Figure 3-14 because their radios can physically transmit at 11 Mb/s. If a higher data rate (such as 18Mb/s) was set to mandatory, only 802.11g clients would be able to associate to the APs.

Setting any of the OFDM rates (rates above 11mb/s) to mandatory disables 802.11b connectivity. This can, for example, allow the administrator to exclude 802.11b clients from the AP by requiring an 802.11g data rate or setting a minimum transmission rate of all clients by disabling 802.11 rates. The reason this might be done is that the same 1500 byte packet at a lower data rate takes a longer time to transmit, and thus, lowers the effective data rate for all wireless clients associated to the AP.

Radio Resource Management (Auto-RF)

In the Cisco WLAN “split MAC” architecture (see Chapter 2, “Cisco Unified Wireless Technology and Architecture,”) the processing of 802.11 data and management protocols and access point capabilities is distributed between a lightweight access point and a centralized WLAN controller. More specifically, time-sensitive activities, such as probe response and MAC layer encryption, are handled at the access point. All other functions are sent to the controller, where system-wide visibility is required.

Real-time RF management of a WLAN network requires system-wide visibility and is implemented at the controller level. The controller learns about the necessary information for an effective RF channel/power plan via information forwarded by the APs in the RF network group.

Note An RF network group (or RF group) is not the same as a mobility group. A mobility group defines a mobility domain of 1–25 controllers in which a client would not be required to change IP address during a roaming event. This is accomplished by building Ethernet over IP tunnels for forwarding client data from an “anchor” controller to the “foreign” controller handling the new AP servicing the client.

Radio Resource Management (RRM), also known as Auto-RF, can adjust the channel (dynamic channel assignment) and power (dynamic transmit power control) to maintain the RF coverage area. It adjusts the power level of the AP to maintain a baseline signal strength with neighboring APs at -65 dBm (configurable) (See Overview of Auto-RF Operation, page 3-29). It adjusts the channel of the AP when it notices nearby interference sources on the channel on which the AP is currently located. It continues to optimize the RF coverage for the best reception and throughput for the wireless network.

RRM understands that the RF environment is not static. As different RF affecting variables change (people in the room, amount of devices stored in the facility, leaves on trees for outside deployment, interference from different RF sources, and so on), the RF coverage adjusts to these variables and changes with them. Because these variables change continuously, monitoring for the RF coverage and adjusting it periodically is necessary.

For more detailed information on Radio Resource Management (Auto-RF), see the following URL: http://www.cisco.com/en/US/products/ps6306/products_white_paper0900aecd802c949b.shtml.

Overview of Auto-RF Operation

Each controller is configured with an RF network group name (called RF Network Name under the WLC Controller -> General menu). In each RF group (if Group Mode is enabled), the controllers elect a leader and form an RF domain. The function of the leader is to collect the network-wide neighbor information from a group of controllers and do the channel/power computation for an optimal system-wide map. If Group Mode is not enabled, the controllers run computations based only on the neighbor data gathered from the APs connected via LWAPP, trying to optimize the signal to -65 dBm between APs.

The APs transmit Radio Resource Management (RRM) neighbor packets at full power at regular intervals. These messages contain a field that is a hash of the RF group name, BSSID, and time stamp. The APs accept only RRM neighbor packets sent with this RF network name.

When neighboring APs receive neighbor messages, they validate them before forwarding them to the controller. If they can validate the message hash and confirm that it belongs to the same RF group, the packet is sent to the controller; otherwise, the AP drops the neighbor packet. The APs then forward the validated messages to the controller, filling in the LWAPP packet status field with the SNR and RSSI of the received neighbor packet.

Table 3-7 provides a summary of the various functions of the devices in the system.

Note TPC performs only downward power level adjustments. Coverage hole detection and correction increases power levels on APs. If a client is associated at a low RSSI level for a significant period of time (and thus has not roamed), it is assumed they cannot find another AP and must be in a coverage hole, and power is increased to help the client communicate.

Auto-RF should not be confused with Rogue Detection (channel scanning), which is done separately from the auto-RF algorithm. APs perform rogue detection by periodically monitoring all country-specific channels (channel scanning). The APs goes “off-channel” for a period not greater than 60 ms to listen to the other channels. Packet headers collected during this time are sent to the controller, where they are analyzed to detect rogue access points, whether service set identifiers (SSIDs) are broadcast or not, rogue clients, ad-hoc clients, and interfering access points.

By default, each access point spends approximately 0.2 percent of its time off-channel. This is statistically distributed across all access points so that no two adjacent access points are scanning at the same time, which can adversely affect WLAN performance. Packets received by the AP from clients are forwarded to the controller with the LWAPP status field filled in, which provides the controller with radio information including RSSI and signal-to-noise ratio (SNR) for all packets received by the AP during reception of the packet.

Table 3-7 Device Function

Device	Functions
RF Group Leader	Collects data from WLCs in the RF group and analyzes it for TX Power Control (TPC) and Dynamic Channel Assignment (DCA) system-wide. TPC adjusts power levels only downward.
Local WLC	Collects data and runs the Coverage Hole Detection and Correction algorithm. Adjusts power levels upward if necessary for clients
Light-weight access point	• Sends neighbor messages on all channels at full power at configured interval • Verifies neighbor hash on received neighbor messages • Scans configured channels for noise, interference, and IDS/rogue detection and alerts if profile fails

Auto-RF Variables and Settings

Auto-RF can be turned on and off via the global setting on the Channel Selection (**Wireless > 802.11b > Configure**) web page (see Figure 3-12). You can manually set the channel and transmit level for the AP from this web page. Additionally, it can be turned off and on from the global Auto-RF web page. Remember that Auto-RF is per band and RF group computations are done for both the 802.11b/g band and another set of computations for 802.11a. The two radios do not have to share to have the same configuration. But these configurations are applied to every AP associated to the controller. Auto-RF configuration variables are shown on the global parameters Auto-RF configuration page (see Figure 3-15).

The first set of variables on the Auto-RF configuration web page corresponds to the RF group. These determine whether the controller joins the dynamic grouping with the other controllers. The dynamic grouping helps the controller find out about APs that are neighbors but might be associated to another controller in the mobility group. If this is disabled, the controller only optimizes the parameters of the access points that it knows about (that is, the ones that are associated to it). The group leader indicates the MAC address of the elected leader. You can find the MAC address of the controller on the inventory web page (you can reach the web page by clicking on **Controller** at the top menu and then **Inventory**).

The Auto-RF configuration web page is divided into three pages, or sections, with a scroll bar that is used to move among the three pages. The first page (see Figure 3-15) is for dynamic channel assignment. This allows the controller to automatically change the channel that the AP is on (for more information, see Dynamic Channel Assignment, page 3-33).

Figure 3-15 Auto-RF (Page 1)

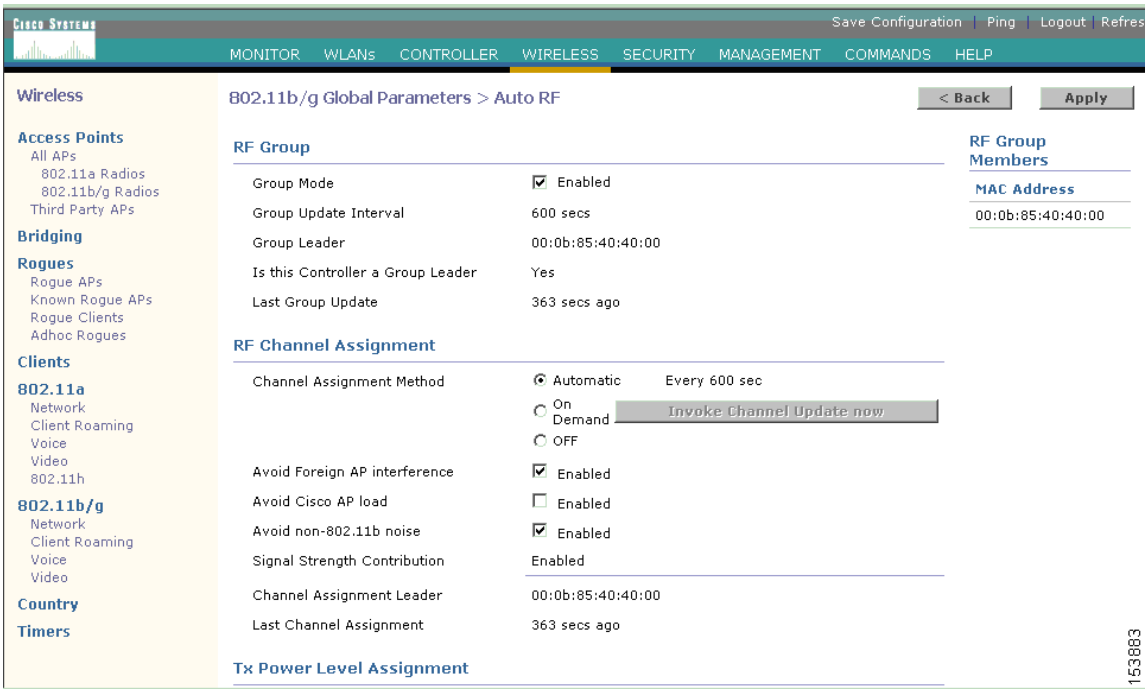

Following the RF channel assignment is the section for assigning the transmit (tx) power level (see Figure 3-16). On this web page, the power level can be fixed for all APs, or it can be automatically adjusted. The web page also indicates the number of neighbors the AP has and the power thresholds for which it is adjusting.

Figure 3-16 Auto-RF (Page 2)

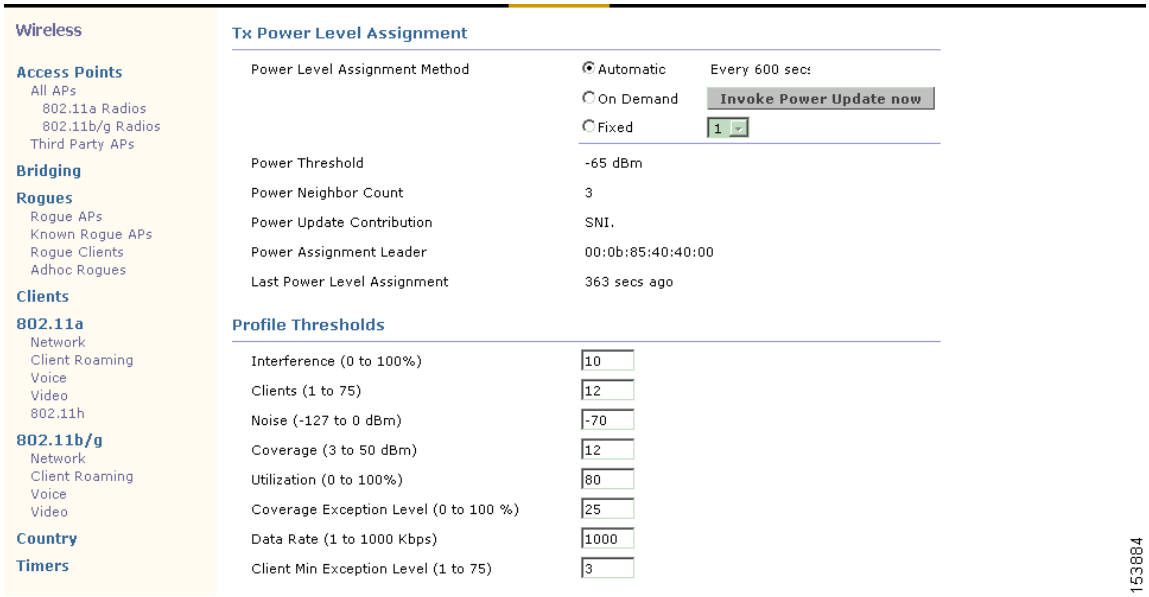

The third web page is for profile thresholds. The controller analyzes the information passed to it by the APs and determines a pass or fail status for each of these thresholds. These pass/fail profiles are best seen in the output of the **show ap auto-rf** *radio ap_name* command (see the following sample). The same information can be seen in graphical form on the **Monitor > 802.11b/g Radios > Detail** web page.

Sample show ap auto-rf Command Output

```
show>ap auto-rf 802.11b <access point name>
Number of Slots . . . . . . . . . . . . . . . . . . . . . . 2
AP Name . . . . . . . . . . . . . . . . . . . . . . . . . . <AP name>
MAC Address . . . . . . . . . . . . . . . . . . . . . . . . 00:0b:85:1b:df:c0
Radio Type . . . . . . . . . . . . . . . . . . . . . . . . . RADIO_TYPE_80211b/g
Noise Information
    Noise Profile . . . . . . . . . . . . . . . . . . . . . PASSED
    Channel 1 . . . . . . . . . . . . . . . . . . . . . . . -93 dBm
    Channel 2 . . . . . . . . . . . . . . . . . . . . . . . -90 dBm
 .
 .
 .
    Channel 11 . . . . . . . . . . . . . . . . . . . . . . -95 dBm
Interference Information
    Interference Profile . . . . . . . . . . . . . . . . . FAILED
    Channel 1 . . . . . . . . . . . . . . . . . . . . . . . -69 dBm @ 31 % busy
    Channel 2 . . . . . . . . . . . . . . . . . . . . . . . -58 dBm @ 26 % busy
 .
 .
 .
    Channel 11. . . . . . . . . . . . . . . . . . . . . . -68 dBm @ 26 % busy
Load Information
    Load Profile . . . . . . . . . . . . . . . . . . . . . PASSED
    Receive Utilization . . . . . . . . . . . . . . . . . . 0 %
    Transmit Utilization . . . . . . . . . . . . . . . . . 0 %
    Channel Utilization . . . . . . . . . . . . . . . . . . 26 %
    Attached Clients . . . . . . . . . . . . . . . . . .  2 clients
Coverage Information
    Coverage Profile . . . . . . . . . . . . . . . . . . .. PASSED
    Failed Clients . . . . . . . . . . . . . . . . . . . . 0 clients
Client Signal Strengths
    RSSI -100 dBm. . . . . . . . . . . . . . . . . . . . . 0 clients
    RSSI -92 dBm . . . . . . . . . . . . . . . . . . . . . 0 clients
 .
 .
 .
    RSSI -52 dBm . . . . . . . . . . . . . . . . . . . . . 1 clients
Client Signal To Noise Ratios
    SNR 0 dBm . . . . . . . . . . . . . . . . . . . . . .  0 clients
    SNR 5 dBm . . . . . . . . . . . . . . . . . . . . . .  0 clients
    SNR 10 dBm . . . . . . . . . . . . . . . . . . . . . 0 clients
 .
 .
 .
    SNR 45 dBm . . . . . . . . . . . . . . . . . . . . . 1 clients
Nearby APs
Radar Information
Channel Assignment Information
    Current Channel Average Energy . . . . . . . . . . . . -68 dBm
    Previous Channel Average Energy . . . . . . . . . . . . -51 dBm
    Channel Change Count . . . . . . . . . . . . . . . . . 21
    Last Channel Change Time . . . . . . . . . . . . . . . Thu Mar 9 12:18:03 2006
    Recommend Best Channel . . . . . . . . . . . . . . . . 11
RF Parameter Recommendations
    Power Level . . . . . . . . . . . . . . . . . . . . . 1
```

```
RTS/CTS Threshold . . . . . . . . . . . . . . . . . . . . . 2347
Fragmentation Threshold . . . . . . . . . . . . . . . . . . 2346
Antenna Pattern . . . . . . . . . . . . . . . . . . . . . . 0
```

The following sections describe some of the Auto-RF variables.

Dynamic Channel Assignment

802.11 MAC uses Carrier-Sense Multiple Access/Collision Avoidance (CSMA/CA). With CSMA/CA, two access points on the same channel (in the same vicinity) get half the capacity of two access points on different channels because of the shared wireless channel. This becomes an issue, for example, when someone reading an e-mail in one business affects the performance of an access point in a neighboring business. Even though these are completely separate networks, someone sending traffic on the common channel can cause network busy signals to clients in the nearby business. This is not to be confused with beacon traffic, which commonly extends well beyond cell range and/or collision domain range. The controllers address this problem and other co-channel interference issues by dynamically allocating access point channel assignments to avoid conflict. Because the controller, or a designated controller called an RF Group Leader, has a system-wide visibility, channels are “reused” to avoid wasting scarce RF resources. In other words, the access point is allocated a different channel, far from the channel used by the neighboring business.

The dynamic channel assignment capabilities of the controller are also useful in minimizing co-channel interference between adjunct access points. For example, with 802.11g, nearby APs cannot both simultaneously use the same channel and receive at 54 Mbps because of interference with each other. By assigning channels, the controller keeps adjacent channels separated, avoiding this problem, as shown in Figure 3-17.

Figure 3-17 Dynamic Channel Assignment

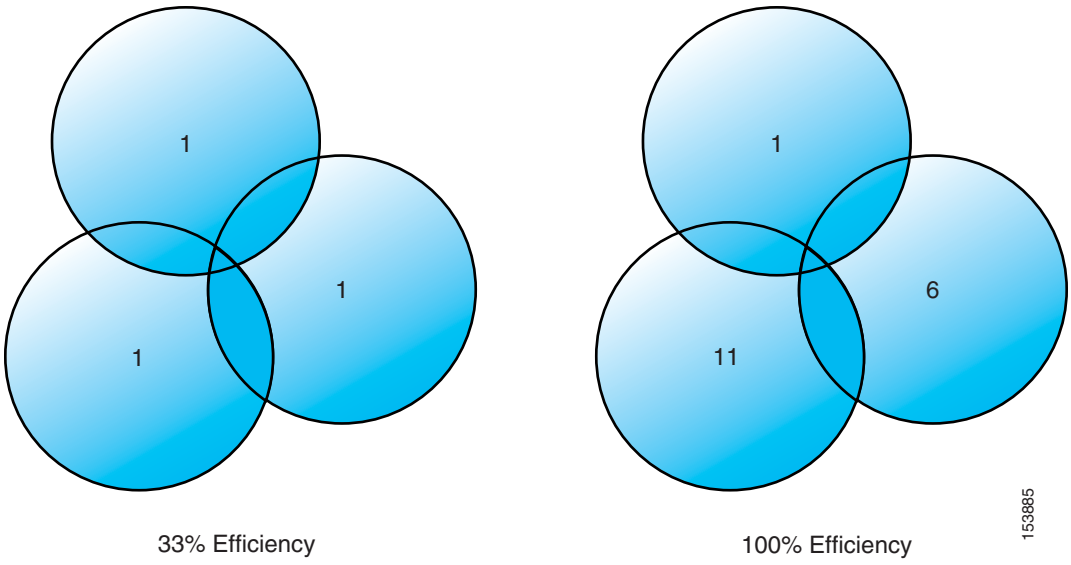

The controller examines a variety of real-time RF characteristics to efficiently handle channel assignments. These include:

- Noise—This limits signal quality at the client and access point, if received in an amount greater than approximately -65 dBm, though this can vary quite easily with directional antennas, range, and periodicity. There are numerous types and effects of interference. An increase in noise reduces the effective cell size at the intervals determined by the controller when it reassesses a radiating

environment. By optimizing channels to avoid noise sources, the controller can optimize coverage while maintaining system capacity. If a channel is unusable because of excessive noise, that channel can be avoided. If other wireless networks are present, the controller shifts the usage of channels to complement the other networks. For example, if one network is on Channel 6, an adjacent WLAN is assigned Channel 1 or 11. This increases the capacity of the network by limiting the sharing of frequencies. If a channel is used so much that no capacity is available, the Cisco Wireless LAN Controller might choose to avoid this channel.

- Client load—Client load is taken into account when changing the channel structure to minimize the impact on the clients currently on the WLAN system. The controller periodically monitors the channel assignment in search of the best assignments. Change occurs only if it significantly improves the performance of the network or corrects the performance of a poorly performing access point.

The controller combines the RF characteristic information to make system-wide decisions. The end result is an optimal channel configuration in a three-dimensional space, where access points on the floor above and below play a major factor in an overall WLAN configuration.

Interference Detection and Avoidance

Interference is defined as any 802.11 traffic that is not part of the Cisco WLAN system; this includes a rogue access point, or a neighboring WLAN. It can also include non-802.11 sources such as certain microwave ovens or cordless phones. It can in certain instances also include various sources of electromagnetic interference (EMI) such as arc welders or federal/military radar facilities. APs are constantly scanning all channels looking for major sources of interference.

If the amount of 802.11 interference hits a predefined threshold, the controller attempts to rearrange channel assignments to optimize system performance in the presence of the interference. This might result in adjacent APs being on the same channel, but logically this is a better choice than staying on a channel that is totally unusable because of an interfering access point.

The controller can respond to a rogue AP on channel 11 by shifting the closest APs to channel 1 or channel 6.

Dynamic Transmit Power Control

The correct AP power settings are essential to maintaining the coverage area, not only to ensure correct (not maximum) amount of power covering an area, but also to ensure that an excess of power is not used, which usually adds significant amounts of noise to the radiating area. AP power settings also control network redundancy by helping to ensure real-time failover in the event of the loss of an AP. The controller is used to dynamically control the AP transmit power level based on real-time WLAN conditions. In normal instances, power can be kept low to gain extra capacity and reduce interference among the APs. RRM attempts to balance access points such that they see their neighbors at -65 dBm. If a failed access point is detected, power can be automatically increased on surrounding access points to fill the coverage gap created by the loss of the AP.

RRM algorithms are designed to create the optimal user experience. For example, if the power of an access point is turned down to Level 4 (where Level 1 = highest and Level 8 = lowest) and the received signal strength indicator (RSSI) value or a user drops below an acceptable threshold, the access point power is increased to provide a better experience to that client.

The change in power settings can be communicated to the clients (via CCX) so that clients see a balanced cell because of the matched transmit power setting and data rate settings between the AP and client. When a client device associates to the access point, the access point sends the maximum power level

setting to the client. If the access points adjust for high data rates and low transmit power but similar adjustments are not made on the clients, the cell might be unbalanced because the client signal can be greater than the access point signal and cause undue interference with neighboring cells.

Note If the access point is configured to disable data rates 1, 2, and 5.5, clients are not required to transmit only at 11 Mbps. In fact, older client devices might not be programmed to recognize the access point configuration for data rates.

Coverage Hole Detection and Correction

If clients on an access point are detected at low RSSI levels, this indicates the existence of an area where clients are continually getting poor signal coverage, without having a viable location to which to roam. The controller adjusts access point power levels to correct the detected hole.

Client and Network Load Balancing

The IEEE standard did not define the process or reasons for client roaming, and therefore it cannot be easily predicted what clients will do in any given situation. For example, all users in a conference room can associate with a single access point because of its close proximity, ignoring other access points that are farther away but much less used.

The controller has a centralized view of client distribution across all access points. This can be used to influence where new clients attach to the network if there are multiple “good” APs available. If configured, the controller can proactively ‘herd’ clients to new access points to improve WLAN performance. This results in a smooth distribution of capacity across an entire wireless network. Keep in mind that this load balancing is done at client association, not while a client is connected.

Cisco Unified Wireless Security

This chapter describes the natively available 802.11 security options and the advanced security features in the Cisco Unified Wireless solution, and how these can be combined to create an optimal WLAN solution.

The Cisco Unified Wireless solution can also be integrated with other Cisco Security solutions; this integration is covered in Chapter 9, "Cisco Unified Wireless Security Integration."

Overview

As network administrators begin to deploy WLANs, they are faced with the challenge of trying to secure these environments while providing maximum flexibility for their users. The Cisco Unified WLAN architecture has multiple components depending on the implementation, but there are two core components that are common in every solution. These are the LWAPP APs -single and dual radio, shown in Figure 4-1, and the Wireless LAN controller (WLC) shown in Figure 4-2.

Figure 4-1 LWAPP APs

Figure 4-2 LWAPP Controller

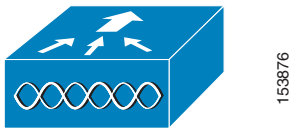

There are various LWAPP AP models and WLC types, but the core WLAN security features remain the same, as does the architecture.

Architecture

The general Cisco Unified WLAN architecture is shown in Figure 4-3, and this architecture can be classified into the following four main layers:

- Client
- Access
- Control and distribution
- Management

Figure 4-3 Unified Wireless Architecture

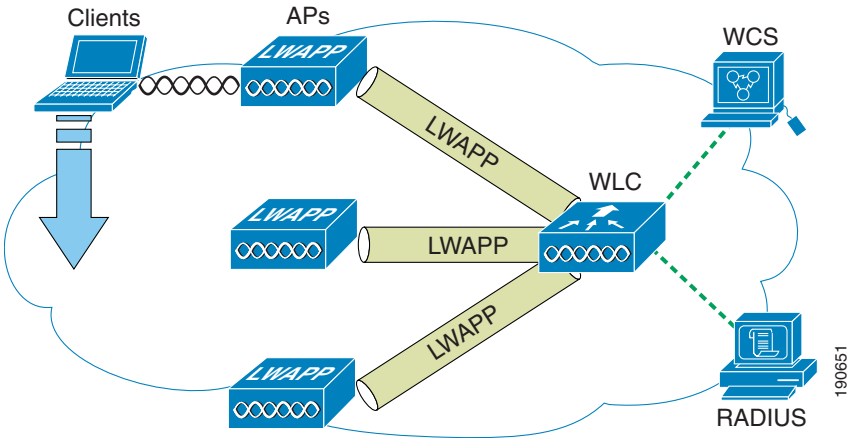

Functional Areas and Components

This section describes the functional areas and components of the Cisco Unified Wireless solution.

Client Component

The client component is critical to the overall security strategy of the solution because the security capabilities of the client often dictate the security capabilities of the solution.

The client device can be a handheld device such as a scanner, PDA, or VoWLAN handset; a mobile device such as a Table PC or laptop computer; or a fixed device such as a PC or printer.

The Cisco Unified Wireless solution is compatible with standard WLAN clients and many specialized WLAN devices. One of the simplest ways to determine which client works best with the Cisco Unified Wireless solution is to consult the Cisco Certified Extensions (CCX) program to verify which WLAN clients are certified for operation with the Cisco solution, in addition to any advanced features included in CCX. For more information on CCX, see the following URL:

http://www.cisco.com/web/partners/pr46/pr147/partners_pgm_concept_home.html

Access Layer

The Access Layer component is the LWAPP APs, which provide the 802.11a/b/g connection for the client devices, and tunnel the client traffic to and from the LWAPP controller across the enterprise network.

Control and Distribution

The Control and Distribution Layer component is primarily performed by the LWAPP controller, which terminates LWAPP tunnels from the LWAPP APs and directs traffic to the appropriate interface and VLAN. The LWAPP controller is also the administrative and authorization interface for APs, and WLAN clients. The LWAPP controller performs additional roles, such as RF management, wireless IDS, and collects location information.

Authentication

A key component in enterprise WLAN deployments is EAP authentication through a RADIUS server. Authentication services for the Cisco Unified Wireless solution can be provided by the Cisco ACS server, which supports all common EAP types including Cisco LEAP, EAP-FAST, EAP-TLS, and PEAP (MSCHAP and GTC), and provides interfaces into external authentication databases such as Microsoft Active Directory, Novell NDS, LDAP, and RSA token servers. The ACS server can also be configured to proxy to other RADIUS servers.

Management

The LWAPP controller has a comprehensive management interface, but centralized management for the Cisco Unified Wireless solution is provided by the Wireless Control System (WCS). In addition to traditional system management functions, WCS provides RF planning and visualization tools, and location services. WCS is covered in more detail later in this document.

WLAN Security Implementation Criteria

For the WLAN network, security is based on both authentication and encryption. Common security mechanisms for WLAN networks are as follows:

- Open Authentication, no encryption
- Wired Equivalent Privacy (WEP)
- Cisco WEP Extensions (CKIP +CMIC)
- Wi-Fi Protected Access (WPA)
- Wi-Fi Protected Access 2 (WPA 2)

WPA and WPA 2 are defined by the Wi-Fi Alliance, which is the global Wi-Fi organization that created the Wi-Fi brand. The Wi-Fi Alliance certifies inter-operability of IEEE 802.11 products and promotes them as the global, wireless LAN standard across all market segments. The Wi-Fi Alliance has instituted a test suite that defines how member products are tested to certify that they are interoperable with other Wi-Fi Certified products.

The original 802.11 security mechanism, WEP, was a static encryption method used for securing wireless networks. Although it applies some level of security, WEP is viewed as insufficient for securing business communications. In short, the WEP standard within 802.11 did not address the issue of how to manage encryption keys. The encryption mechanism itself was found to be flawed, in that a WEP key could be derived simply by monitoring client traffic. Cisco WLAN products addressed these issues by introducing 802.1x authentication and dynamic key generation and by introducing enhancements to WEP encryption: CKIP and CMIC. 802.11i is a standard introduced by the IEEE to address the security shortcomings of the original 802.11 standard. The time between the original 802.11 standard and the ratification of 802.11i saw the introduction of interim solutions.

WPA is an 802.11i-based security solution from the Wi-Fi Alliance that addresses the vulnerabilities of WEP. WPA uses Temporal Key Integrity Protocol (TKIP) for encryption and dynamic encryption key generation by using either a pre-shared key, or RADIUS/802.1x-based authentication. The mechanisms introduced into WPA were designed to address the weakness of the WEP solution without requiring hardware upgrades. WPA2 is the next generation of Wi-Fi security and is also based on the 802.11i standard. It is the approved Wi-Fi Alliance interoperable implementation of the ratified IEEE 802.11i standard. WPA 2 offers two classes of certification: Enterprise and Personal. Enterprise requires support for RADIUS/802.1x-based authentication and pre-shared key (Personal) only requires a common key shared by the client and the AP. The new AES encryption mechanism introduced in WPA2 generally requires a hardware upgrade from earlier versions of WLAN clients and APs, however all Cisco LWAPP APs support WPA2.

Table 4-1 summarizes the various specifications.

Table 4-1 WLAN Security Mechanisms

Feature	Static WEP	802.1x WEP	WPA	WPA 2 (Enterprise)
Identity	User, machine or WLAN card	User or machine	User or machine	User or machine
Authentication	Shared key	EAP	EAP or pre-shared keys	EAP or pre-shared keys
Integrity	32-bit Integrity Check Value (ICV)	32-bit ICV	64-bit Message Integrity Code (MIC)	CRT/CBC-MAC (Counter mode Cipher Block Chaining Auth Code - CCM)
Encryption	Static keys	Session keys	Per Packet Key rotation via TKIP	CCMP (AES)
Key distribution	One time, Manual	Segment of PMK	Derived from PMK	Derived from PMK
Initialization vector	Plain text, 24-bits	Plain text, 24-bits	Extended IV-65-bits with selection/sequencing	48-bit Packet Number (PN)
Algorithm	RC4	RC4	RC4	AES
Key strength	64/128-bit	64/128-bit	128-bit	128-bit
Supporting infrastructure	None	RADIUS	RADIUS	RADIUS

The Cisco Wireless Security suite provides the user with the options to provide varying security approaches based on the required or pre-existing authentication, privacy and client infrastructure. Cisco Wireless Security Suite supports WPA and WPA2, including:

- Authentication based on 802.1X using the following EAP methods:
  - Cisco LEAP, EAP-Flexible Authentication via Secure Tunneling (EAP-FAST)
  - PEAP- Generic Token Card (PEAP-GTC)
  - PEAP-Microsoft Challenge Authentication Protocol Version 2 (PEAP-MSCHAPv2)
  - EAP-Transport Layer Security (EAP-TLS)
  - EAP-Subscriber Identity Module (EAP-SIM)
- Encryption:

- AES-CCMP encryption (WPA2)
- TKIP encryption enhancements: key hashing (per-packet keying), message integrity check (MIC) and broadcast key rotation via WPA TKIP Cisco Key Integrity Protocol (CKIP) and Cisco Message Integrity Check (CMIC)
- Support for static and dynamic IEEE 802.11 WEP keys of 40 bits, 104, and 128 bits

Note 128 bit WEP (128 bit WEP key =152 bit total key size as IV is added to key) is not supported by all APs and clients. Even if it was, increasing WEP key length does address the inherit security weaknesses of WEP.

IPsec

In addition to the variety of security mechanism supported natively in 802.11, authentication and encryption can also be performed at higher network layers. The most common mechanism being IPsec, which is typically implemented in place of or in addition to 802.11 security mechanisms.

The operation of IPsec is not covered in this chapter; however, where appropriate, IPsec-related features and design recommendations for WLAN deployments are made.

802.1x/EAP Authentication

802.11i specifies the use of 802.1x for providing port access control on WLAN network ports. WPA, and WPA2 further specify the use Extensible Authentication Protocol (EAP) to exchange authentication information. EAP payloads are placed within 802.1x frames or RADIUS packets to establish communication between the supplicant -WLAN client, and the Authenticator = AP/WLC -RADIUS server. Access to the network is determined by the success or failure of the EAP authentication, and the WLAN encryption is derived from shared cryptographic data created during the EAP authentication. Figure 4-4 shows the general authentication flow.

Figure 4-4 Generic EAP over 802.1x Authentication Mode

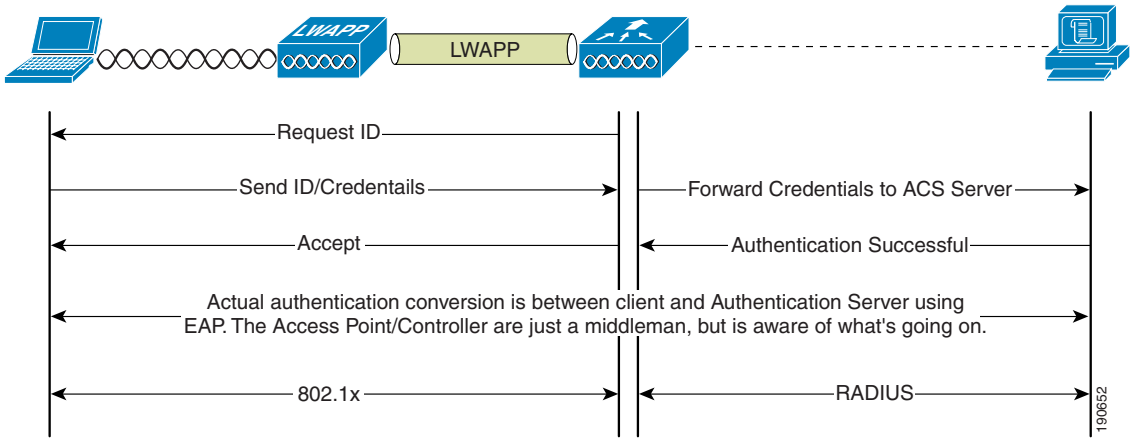

Various EAP types are used in WLAN solutions. Some common EAP types are the following:

- EAP-TLS (transport layer security-PKI-based client and server authentication)

- Cisco Lightweight Extensible Authentication Protocol (LEAP)
- Protected Extensible Authentication Protocol (PEAP)
- Flexible Authentication via Secured Tunnel (EAP-FAST)

These EAP types define how the authentication messaging takes place between the client and the authentication server. The Supplicant and the Authentication Server must support the same EAP types. Because the EAP payloads are passed across the Authenticator without being parsed, the Authenticator need not care about the EAP authentication type. EAP payload data of interest to the Authenticator comes from a successful authentication. Such data might include RADIUS VSAs specifying the VLAN ID to be used by the client, ACLs, or controlling QoS parameters.

Although the Authenticator need not know the EAP type used, Authenticator configuration can impact the successful implementation of a given EAP type; for example, the 802.1x timeouts and retries parameters can impact the usability of PEAP-GTC because it requires a user to enter data.

Table 4-2 provides a brief comparison of various EAP supplicants.

Table 4-2 EAP Authentication Comparison

	Cisco LEAP	Cisco EAP-FAST	PEAP/MS-CHAPv2	PEAP(EAP-GTC)	EAP-TLS
Single sign-on (MSFT AD only)	Yes	Yes	Yes	Yes[1]	Yes
Login scripts execution (MSFT AD only)	Yes	Yes	Yes	Some	Yes[2]
Password Change (MSFT AD)	No	Yes	Yes	Yes	N/A
Cisco 350 and CB20A client support for Windows XP, 2000, and Windows CE OS	Yes	Yes	Yes	Yes	Yes
PCI card client support for Windows XP and Windows 2000	Yes	Yes	Yes	Yes	Yes
Microsoft AD DB support	Yes	Yes	Yes	Yes	Yes
ACS local DB support	Yes	Yes	Yes	Yes	Yes
LDAP DB support	No	Yes[3]	No	Yes	Yes
OTP authentication support	No	Yes [8]	No	Yes	No
RADIUS server certificate required?	No	No	Yes	Yes	Yes
Client certificate required?	No	No	No	No	Yes
Susceptible to Dictionary attacks?	Yes[4]	No	No	No	No
Susceptible to MITM attacks?	No	No[5]	Yes[6]	Yes[7]	No
Fast secure roaming (Cisco CCKM)	Yes	Yes	Yes[1]	Yes[1]	Yes[1]
Local authentication	Yes	Yes	No	No	No
WPA support (Windows 2K/XP)	Yes	Yes	Yes	Yes	Yes
Proactive Key Caching (PKC WPA2 802.11i Fast Roaming)	Yes	Yes	Yes	Yes	Yes

[1] Supplicant Dependent

[2] Machine account on Windows AD is required to enable Login Script execution for PEAP and EAP-TLS

[3] Automatic provisioning is not supported for LDAP back-end DBs. Manual provisioning would have to be used for back-end LDAP DBs.

[4] Strong Password policy is required for LEAP deployment to mitigate risks because of offline (such as passive) dictionary attacks.

[5] EAP-FAST with automatic provisioning is susceptible to rogue server (reduced MITM) attack during the phase 0 (automatic provisioning stage). MITM attacks require the attacker to spoof a legitimate AP. Which means strategies such as Rogue AP detection and Management Frame Protection can detect the presence of these attacks.

[6] PEAP (specifically PEAPv1) is vulnerable to MITM attacks. This is covered in a document at the following URL:

http://www.ietf.org/internet-drafts/draft-puthenkulam-eap-binding-04.txt.

This MITM vulnerability will be fixed in PEAPv2.

[7] Although Cisco PEAP, as a hybrid authentication type, is theoretically vulnerable to MITM attacks, the Cisco supplicant implementation of PEAPGTC is less vulnerable, as it does not accept the same authentication types inside and outside the TLS tunnel, a requirement for the MiTM exploit publicly detailed. OTP Authentication supported in EAP-FAST v1a.

[8] For comment on EAP-FAST OTP support Supplicant Dependent

Wired Equivalent Privacy

This section provides a brief description of encryption and message integrity mechanisms (see Figure 4-5). The main goals for encryption and message integrity are to prevent disclosure, modification, and insertion of packets in a WLAN.

References to sources that provide more detailed information and an analysis of crypto-algorithms, key management, and implementations can be found in References, page 4-11.

Figure 4-5 WEP Encapsulation Process

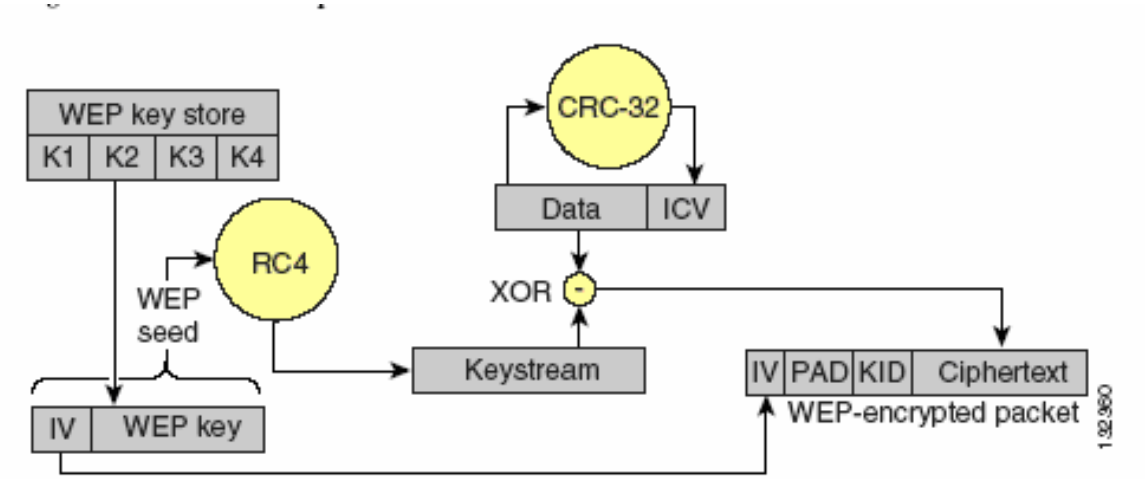

The LWAPP WLAN solution supports three key lengths: the standard 40 and 104 bit key lengths, and an additional 128 bit key. The use of the 128 bit key is not recommended because 128 bit keys are not widely supported in WLAN clients, and the additional key length does not address the weakness inherent in WEP encryption

Temporal Key Integrity Protocol

With TKIP, the main objective is to address the problems with WEP and to work with legacy hardware; therefore, the base encryption mechanism is still RC4, the same as WEP.

TKIP is a cipher suite that includes key mixing algorithms and a packet counter to protect the keys. It also includes the Michael Message Integrity Check (MIC) algorithm that, along with the packet counter, can prevent packet modification and insertion. Figure 4-6 illustrates the TKIP encapsulation process.

Figure 4-6 TKIP Encapsulation Process

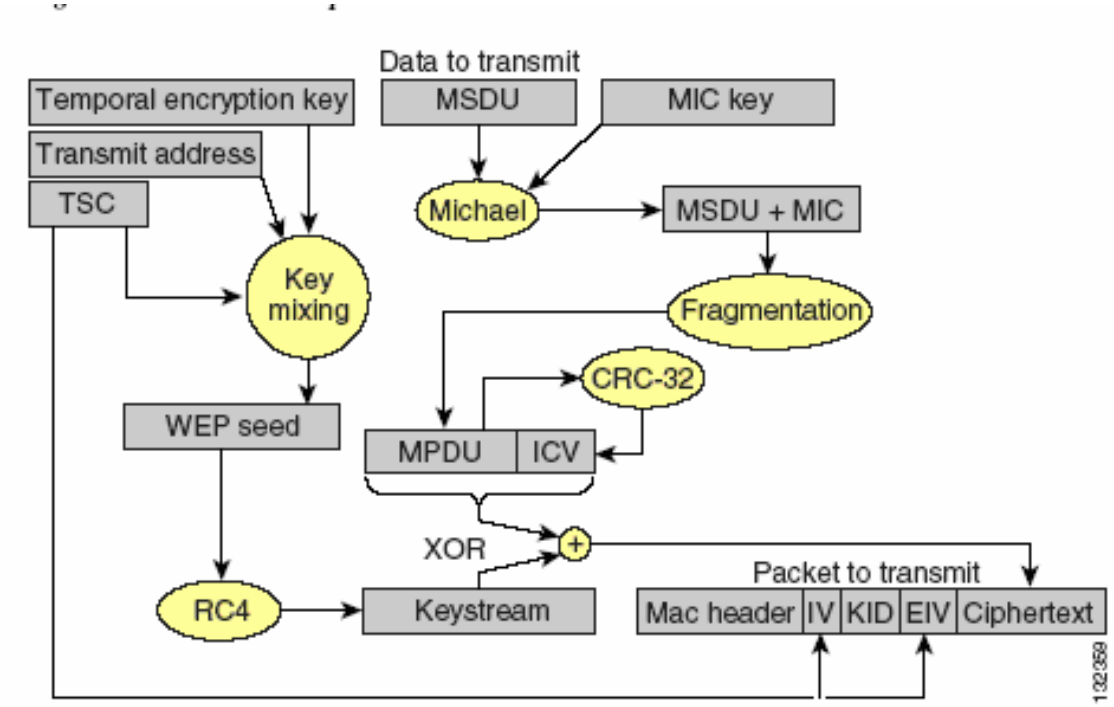

Cisco Key Integrity Protocol and Cisco Message Integrity Check

Cisco Key Integrity Protocol (CKIP) and Cisco Message Integrity Check (CMIC) are the Cisco versions of TKIP and MIC, respectively. CKIP and CMIC were developed to address the WEP vulnerabilities before the release of WPA. Combined, CKIP and CMIC provide encryption and message integrity far superior to WEP.

Counter Mode/CBC-MAC Protocol

Counter Mode/CBC-MAC Protocol (CCMP) is an algorithm based on the Advanced Encryption Standard (AES). It provides encryption and data integrity, and is part of the 802.11i specification. AES has stronger encryption and message integrity than TKIP, but is not compatible with legacy WLAN hardware because of the much more intensive processing required for AES encryption and decryption. Figure 4-7 illustrates the CCMP encapsulation process.

Figure 4-7 CCMP Encapsulation Process

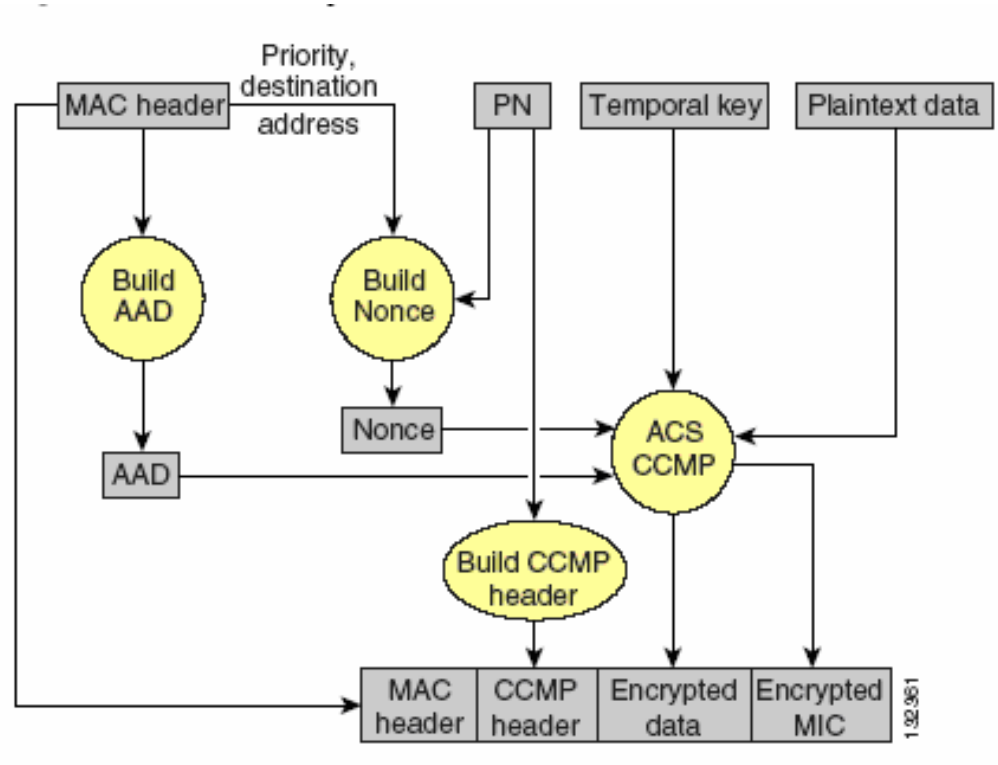

Proactive Key Caching and CCKM

Proactive Key Caching (PKC) is an 802.11i extension that allows for the proactive caching (before the client roaming event) of the Pair-wise Master Key (PMK) that is derived during a client 802.1 x/EAP authentication at the AP (see Figure 4-8). If a PMK (for a given WLAN client) is already present at an AP when presented by the associating client, full 802.1x/EAP authentication is not required. Instead, the WLAN client can simply use the WPA four-way handshake process to securely derive a new session encryption key for communication with that AP.

The distribution of these cached PMKs to APs is greatly simplified in the Unified Wireless deployment. The PMK is simply cached in the controller(s) and made available to all APs that connect to that controller, and between all controllers that belong to the mobility group of that controller in advance of a client roaming event.

Figure 4-8 Proactive Key Caching Architecture

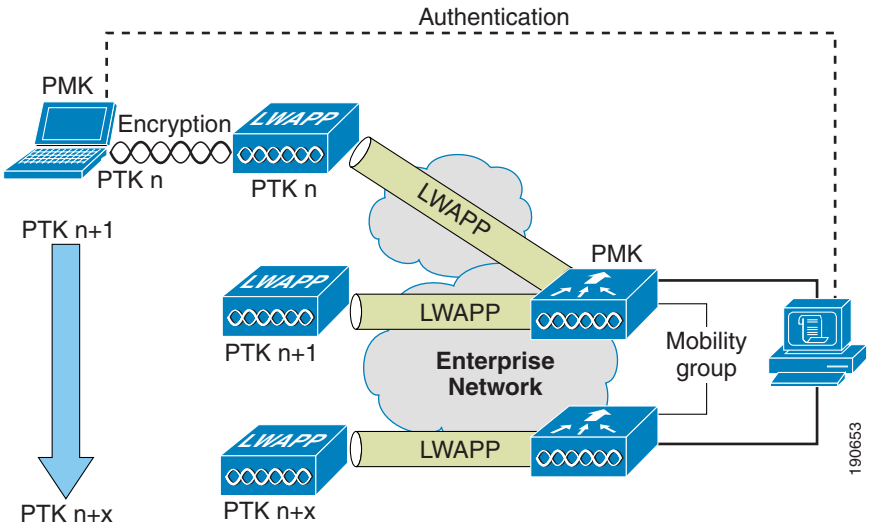

Cisco Centralized Key Management (CCKM) is a Cisco standard supported by CCX clients to provide Fast Secure Roaming. The principle mechanism for accelerating roaming is the same as PKC, by using a cached PMK, but the implementation is slightly different and the two mechanisms are not compatible. A detailed description of FSR and CCKM can be found at the following URL:

http://www.cisco.com/en/US/products/hw/wireless/ps430/prod_technical_reference09186a00801c5223.html

The state of the each WLAN client's key caching can be seen with the **show pmk-cache all** command This identifies which clients are caching the keys, and which key caching mechanism is being used.

The 802.11r workgroup is responsible for the standardization of a fast secure roaming mechanism for 802.11. The WLC controller supports both CCKM and PKC on the same WLAN -802.1x+CCKM, as shown in the following example:

```
WLAN Identifier.................................. 1
Network Name (SSID).............................. wpa2
…
Security
   802.11 Authentication:........................ Open System
   Static WEP Keys.............................. Disabled
   802.1X....................................... Disabled
   Wi-Fi Protected Access (WPA/WPA2)............. Enabled
      WPA (SSN IE)............................... Disabled
      WPA2 (RSN IE).............................. Enabled
         TKIP Cipher............................. Disabled
         AES Cipher.............................. Enabled
   Auth Key Management
         802.1x.................................. Enabled
         PSK..................................... Disabled
         CCKM.................................... Enabled
…
(Cisco Controller) >show pmk-cache all
PMK-CCKM Cache
                              Entry
Type        Station           Lifetime   VLAN Override        IP Override
------    --------------    --------   ------------------   ---------------
CCKM    00:12:f0:7c:a3:47   43150                             0.0.0.0
RSN    00:13:ce:89:da:8f    42000                             0.0.0.0
```

References

There are many articles and books that cover security in detail, such as the following:

- *Cisco Wireless LAN Security* by Sankar, Sundaralingam, Balinsky and Miller
- *802.11 Real Security* by Edney and Arbaugh
- *802.11 Wireless Fundamentals* by Roshan and Leary

WLAN Security Selection

There are many options for selecting and implementing the security standards for WLANs. However, in most implementations, the decisions are bound by existing enterprise security practices and clients participating in the WLANs, When dealing with clients, you need to know what supplicants are available for those clients, and specifically what authentication/identity framework is used by the enterprise.

Given these options, the decision of what must be implemented can be varied and challenging. Cisco provides the ability to segment various security schemes via VLANs, which is described in a separate white paper.

The following tables compare and summarize the security standards for WLANs. Table 4-3 compares Cisco LEAP, PEAP, and EAP-TLS.

Table 4-3 Comparing LEAP, PEAP, EAP-TLS

Cisco LEAP	Supports many operating systems (Windows 95, 98, 2000, XP, Me, NT, Mac OS, Linux, DOS, Windows CE)
	Supports many adapters and client devices, including devices with small processors
	Supports a variety of wireless LAN devices like Cisco workgroup bridges, wireless bridges, and repeaters
	Does not require certificates or a Certificate Authority
	Can be configured quickly and easily
	Supports a single sign-on with an existing user name and password
	Has been field-proven since 2001
	Requires minimal client software overhead
	Utilizes minimal authentication messaging
	Known security exposure—requires strong passwords
EAP-FAST	Tunnel establishment is based on shared secret keys that are unique to users. (Protected Access Credentials (PACs) and can be distributed automatically (Automatic or In-band Provisioning) or manually (Manual or Out-of-band Provisioning) to client devices.)
	Single sign-on (SSO) using the user name and password supplied for Windows networking logon

Table 4-3 Comparing LEAP, PEAP, EAP-TLS (continued)

	Wi-Fi Protected Access (WPA) support without third-party supplicant (Windows 2000 and XP only)
	Support for key Cisco Unified WLAN Architecture features: Fast Secure Roaming (CCKM) and Local
	RADIUS Authentication
	No reliance on Microsoft 802.1X framework
	No certificates authority needed/ No requirement for certificates
	Windows Password Aging (support for server-based password expiration)
EAP-TLS	Supported natively on Windows XP and Windows 2000 (with service pack)
	Supports NDS and LDAP (when appropriately configured)
	Uses same PKI mechanism as wired or dial-up access for easy distribution of client certificates
	Official EAP type tested with Wi-Fi Protected Access (WPA)– although other EAP types will work with WPA
	Exposes user information in the certificate
PEAP-MSCHAP	Supports password change at expiration
	Is defined in a draft RFC
	Does not expose the logon user name in the EAP Identity Response
	Is not vulnerable to a dictionary attack
	Requires a server certificate and CA certificate, but does not require per-user certificates
	The authentication protocol is protected by a TLS tunnel but the tunneled authentication protocol is limited to MSCHAPv2
	Supported natively on Windows XP and Windows 2000(with service packs),
	Integrates into Active Directory user database
PEAP-MSCHAPv2	Support for key Cisco Unified WLAN Architecture features: Fast Secure Roaming (CCKM) and Local
	RADIUS Authentication
	No reliance on Microsoft 802.1X framework
	No certificates authority needed/ No requirement for certificates
PEAP-GTC	Supports authentication using one-time passwords
	Supports NDS and LDAP
	Supports password change at expiration
	Is defined in a draft RFC
	Does not expose the logon user name in the EAP identity response
	Is not vulnerable to a dictionary attack
	Requires a server certificate and CA certificate, but does not require per-user certificates

Table 4-4 lists the advantages of using 802.1x EAP for WLAN.

Table 4-4 802.1x Comparison to IPsec VPN

802.1x EAP Types versus IPsecVPNs	The advantages of using 802.1X EAP for WLAN are:·
	Included with Wi-Fi certified clients and access points
	Minimal client software overhead
	Minimal authentication messaging overhead
	Minimal management overhead
	Natively supported on many operating systems
	Layer 3 roaming support
	Authentication choice for enterprise deployments

Table 4-5 compares the advantages of Cisco TKIP with WPA TKIP.

Table 4-5 Cisco KIP Comparison to WPA TKIP

Cisco TKIP	WPA TKIP
Cisco TKIP is well-suited to the following deployments: • Enhanced security is required but a WPA supplicant cannot be supported on the client platform. • If 802.1q trunks are supported by the Layer 2 infrastructure and it is possible to use WLAN VLANs to segregate Cisco TKIP users from other WLAN users.	WPA TKIP is well suited to the following deployments: • Client devices can support WPA. • Cisco Compatible version 2 cards in use. • If 802.1q trunks are not supported by the Layer 2 infrastructure WPA and non-WPA clients can operate on the same SSID, via WPA migration mode. Native support for wireless devices and authentication protocol is desired (no external supplicant required).

Table 4-6 lists the advantages and disadvantages of using VPN for WLAN.

Table 4-6 Advantages and Disadvantages of Using VPN for WLAN

Advantages	Disadvantages
Uses 3DES or AES encryption	Client software overhead
Enforces remote user authentication and polices for Wireless LAN users	Authentication messaging overhead
Leverages existing VPN if already installed for wired network	Management overhead because one VPN application is required per client
Used for remote users accessing the network while on the road at airports, hotels, conference centers	Does not support single sign on using Windows log-in
	Client traffic is hidden from WLAN infrastructure, limiting the application of any policies based on client traffic
	Limited or no multicast and multiprotocol support

WLAN Security Configuration

The WLC allows the configuration of multiple WLANs that can be mapped to different dot1q interfaces on the WLC, and the WLANs can applied to different APs through AP grouping.

Figure 4-9 shows the main configuration page for WLAN security on WLC. This is part of the WLAN menu; each WLAN that is created has a similar page where key 802.11 parameters can be configured, as well as the security settings for that WLAN. These security settings include the type of authentication and encryption to be used for that WLAN, including any sub-options applicable to that security option. For example, solutions that require 802.1x based authentication allow RADIUS servers to be selected for that authentication type.

Figure 4-9 WLAN Configuration Page

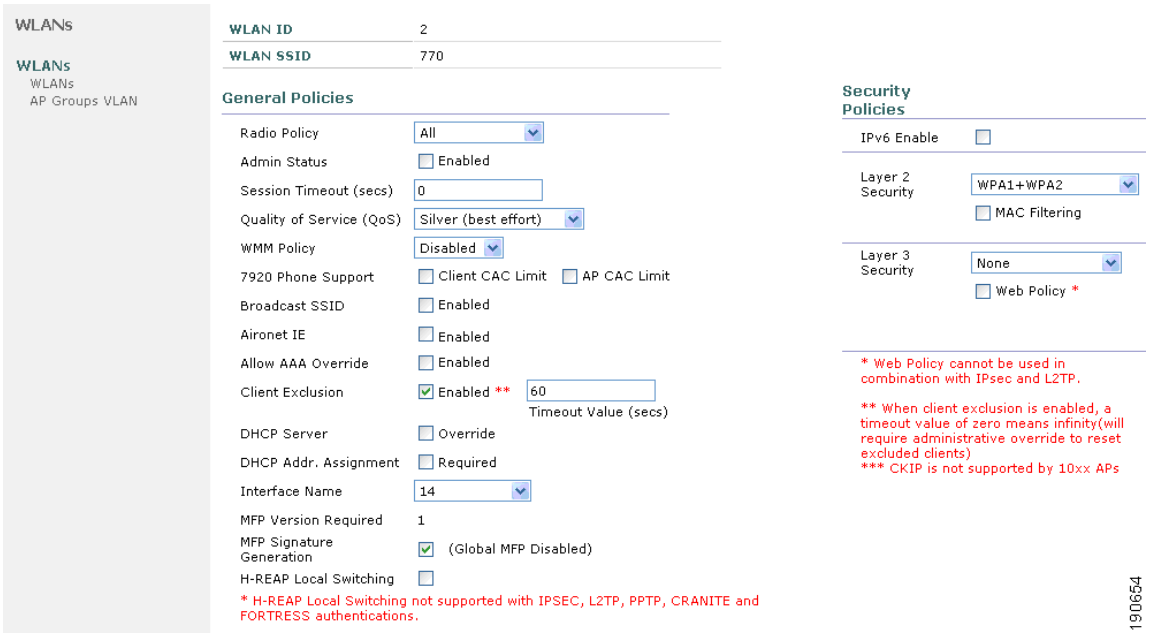

Figure 4-10 shows the various Layer 2 security options that are available on the WLAN. These range from Open Authentication with no encryption to WPA-2.

Figure 4-10 Controller WLAN Layer 2 Security Options

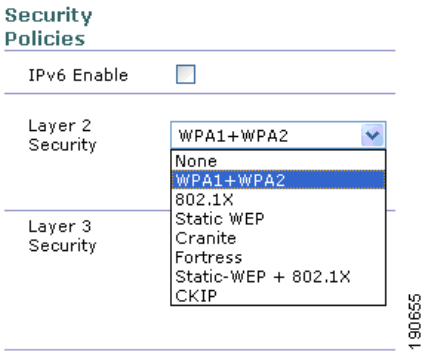

The RADIUS servers used in the WLAN configuration are configured on the controller in the security section, shown in Figure 4-11. Multiple RADIUS servers can be configured, and assigned different priorities. Note that the RADIUS server priority setting from Figure 4-11 is not the priority of the RADIUS servers used in the WLAN authentication, that priority is established on the WLAN configuration page.

The Retransmission timeout sets the delay between retransmission if the RADIUS server does not respond to the RADIUS request. The WLC retries five times before trying the next RADIUS server in a configured list.

Note that the WLC does not automatically retry the preferred RADIUS server when it has failed over to another server, unless that server stops responding; for example, the RADIUS server does not fail back.

Note also that the source address used by the controller for AAA authentication is the management address of the WLC.

Figure 4-11 RADIUS Configuration

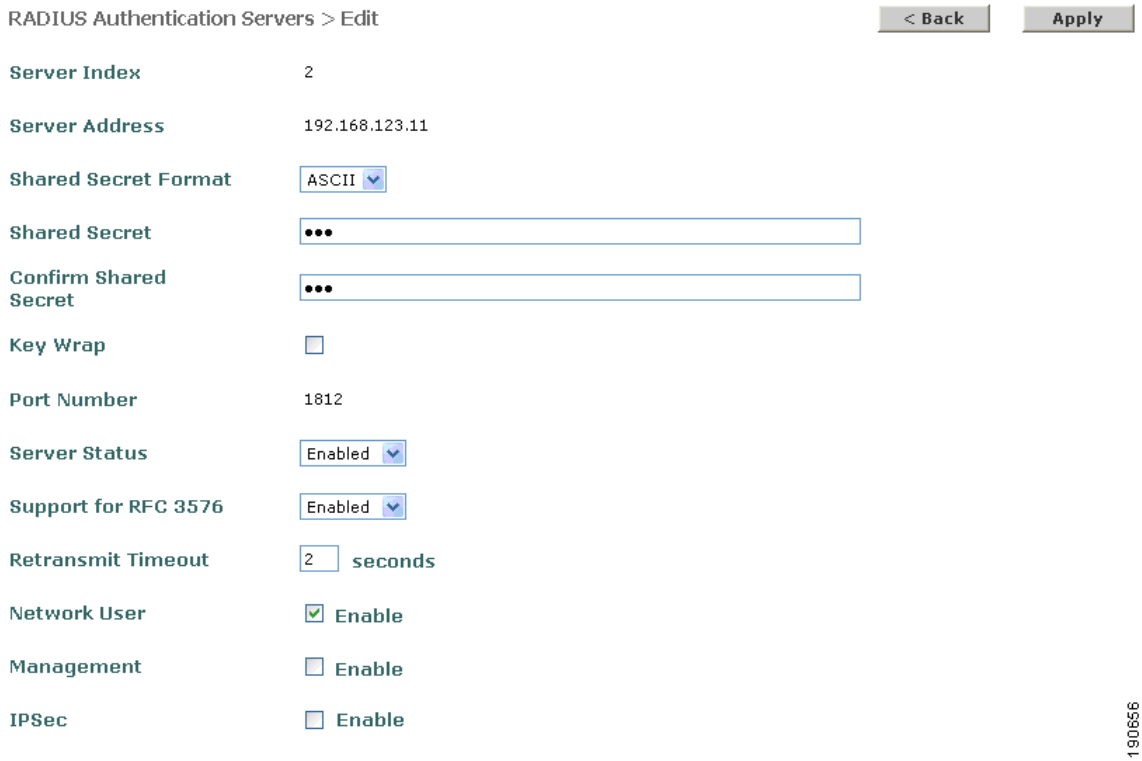

The Key WRAP option should be left unchecked unless a RADIUS server using the Key WRAP features (typically in a FIPS compliant implementation) is being configured.

Unified Wireless Security

The Cisco Unified WLAN Architecture addresses many facets of WLAN security, and although this white paper focuses on WLAN Data Transport Security, a brief description of the other security features of the solution is described in this section. The security features are grouped into the following three categories:

- Infrastructure Security—Security features addressing the configuration and deployment of the WLAN solution itself
- WLAN Data Transport Security—The security features addressing the WLAN traffic
- WLAN Environment Security—The security features designed to protect the WLAN environment and resources from attack or accidental interference

Infrastructure Security

The deployment of WLANs in enterprises generally involves the deployment of enterprise network equipment in locations other than locked wiring closets, or Network Operating Centers (NOC). This introduces a new exposure to some networks, because it increases the likelihood of the theft or attacks on network equipment, which can in turn expose authentication keys, encryption keys, passwords, and other configuration data relating to network security.

The Cisco Unified WLAN Architecture is immune to the vulnerabilities described above by virtue of the fact that the centralized architecture does not store any security configuration information in NVRAM within the LWAPP APs themselves (configuration is lost when power is removed from the AP). Instead, all configurations related to WLAN and system security are implemented in the LWAPP controller, which is typically deployed in a secured location. The privacy of network configuration is further enhanced by its encryption between the LWAPP AP and the LWAPP controller, and by preventing console access to the LWAPP AP configuration. This prevent the WLAN configuration information being learned through capturing the LWAPP stream of reading the configuration on an active AP.

The Cisco Unified WLAN Architecture also prevents the threat of impersonation and spoofing to gain access to network configuration information through the use of X.509 certificates on the LWAPP devices, and also requires PKI authentication before LWAPP configuration information is exchanged. In addition, the MAC addresses contained in the X.509 certificates can be authenticated against a centralized database(s) to ensure that unauthorized APs do not connect to a controller.

The WLC, which is the core component of the Cisco Unified WLAN solution, uses dot1q VLANs to provide isolation between user WLAN traffic, and the WLC's management interfaces. The WLC also offers secure management access using SSH, HTTPS, and SNMPv3 protocols, as well as providing an out of band management interface on many WLC models.

Additionally, the WLC allows ACLs to be implemented to further restrict access. This is accomplished by using the **config acl cpu** command. Applying ACLs directly to the Management and AP-Management WLC interfaces currently has no effect on traffic to the WLC, and only applies to WLAN client traffic on those interfaces. Therefore, when using ACLs to control traffic to the WLC management interfaces, use the **config acl cpu** command.

WLAN Data Transport Security

The Cisco Unified WLAN Architecture provides a full range of WLAN transport security features ranging from open unauthenticated connections to WPA2 connections. These various security models can be supported over the same infrastructure, and mapped to different wired network connections through configuration policies supplied by the controller or from a AAA server.

The Cisco Unified WLAN Architecture also resolves the architectural challenge of segmenting WLAN traffic from wired data traffic by using LWAPP tunnels to transport WLAN user and control data between APs and the controller and then uses other LWAPP controller features such as 802.1q VLANs and or EoIP tunnels to provide further segmentation.

WLAN Environment Security

The Cisco Unified WLAN Architecture users RF Security features to detect and avoid 802.11 interference and control unwanted RF propagation. The WLAN Intrusion Prevention and Location features not only detect rogue devices or potential WLAN threats, but also locates these devices. This enables system administrators to quickly assess the threat level and take immediate action to mitigate threats as required.

A key component that facilitates WLAN Environment Security reporting is the WCS server. WCS collects and correlates information from the WLCs, and links this information with preconfigured location information stored in the WCS.

The WCS is described in more detail in a subsequent chapter.

Rogue AP

A standard AP looks for rogue activity by going off channel for 50 ms to listen for rogue APs, clients, monitor for noise, and channel interference. The channels to be scanned are configured in the global WLAN network parameters for 802.11a and 802.11b/g. Any detected rogue clients or APs are sent to the controller, which gathers the following:

- Rogue AP MAC address
- Rogue AP name
- Rogue connected clients' MAC address
- Whether the frames are protected with WPA or WEP
- The preamble
- SNR
- RSSI

The WLC waits to label this as a rogue client or rogue AP because it might not have been reported by another AP until it completes another scanning cycle (the WLC ensures that its AP and client database are up to date before labeling a client or AP as rogue). The same AP again moves to the same channel to monitor for rogues access points/clients, noise and interference. If the same clients and/or access points are detected, they are listed as a rogue on the controller again. The WLC now begins to determine whether this rogue is attached to the local network or simply a neighboring AP. In either case, an AP that is not part of the managed local WLAN is considered a rogue.

If an AP is configured for "monitor mode", it does not carry user traffic but spends all its time scanning different channels.

If an AP is configured as a rogue detector, its radio is turned off and its role is to listen for MAC addresses, detected by the controller as being rogue APs. The rogue detector listens for ARP packets, and to be effective should be connected to all broadcast domains via trunk link if desired to maximize the likelihood of detection; the AP is still connected to the network via the native VLAN, but monitors other VLANs for ARP frames.

Rogue detector APs might not be practical for some deployments, and do not discover clients that are going through a WLAN router, which are common consumer devices. Rogue Location Discovery Protocol can aid in these cases, where a standard AP, on detecting a rogue AP, can attempt to associate with the rogue AP as a client and send a test packet to the controller. This confirms that the rogue AP is actually on the network The IP addressing information obtained from the test packet can be used to determine the location of the rogue on the network.

Management Frame Protection

One of the challenges in 802.11 has been that management frames are sent unprotected, and are therefore vulnerable to spoofing attacks. To address this, Cisco has created a digital signature mechanism to insert a Message Integrity Check into the 802.11 management frames (see Figure 4-12). This allows legitimate members of a WLAN deployment to be identified, and facilitates easy detection of rogue infrastructure devices through the absence of valid MICs in their management frames.

The message integrity check that is used in MFP is not a simple CRC hashing of the message, but also includes a digital signature component. The MIC component of MFP ensures that a frame has not been tampered with, and the digital signature component ensures that the MIC can have only been produced by a valid member of the WLAN domain. The digital signature key used in MFP is shared among all controllers in a mobility group; different mobility groups have different keys. This allows the validation of all WLAN management frames processed by the WLCs in that mobility group.

Figure 4-12 Management Frame Protection

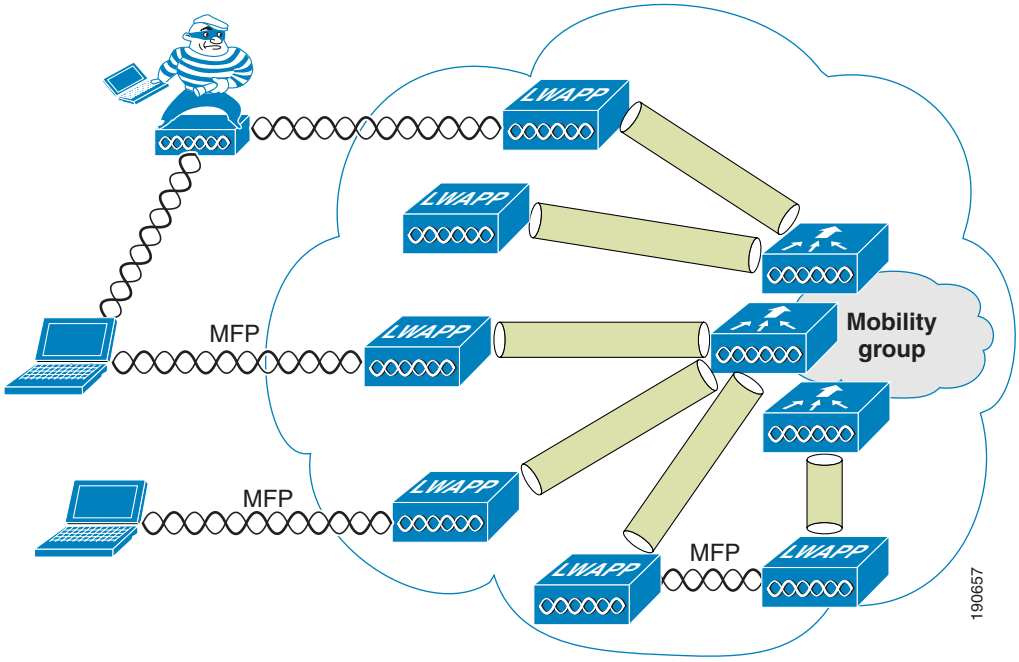

Currently, MFP is only possible for WLAN infrastructure, but with CCX v5, WLAN clients will be able to learn the mobility group MFP key, and therefore detect and reject invalid frames.

Management Frame Protection provides the following benefits:

- Provides for the authentication of 802.11 management frames by the WLAN network infrastructure
- Allows detection of malicious rogues that are spoofing a valid AP MAC or SSID to avoid detection as a rogue AP, or as part of a man-in-the-middle attack
- Increases the quality of rogue AP and WLAN IDS signature detection
- Will provide protection of client devices with CCX v5
- Also supported with Autonomous AP/ WDS/ WLSE in version 12.3(8)/ v2.13

There are two steps to enable MFP; one to enable it on the WLC, and the second to enable it on the WLAN that is part of the mobility group. Figure 4-13 shows the enabling of MFP on the WLC.

Figure 4-13 Enabling MFP on the Controller

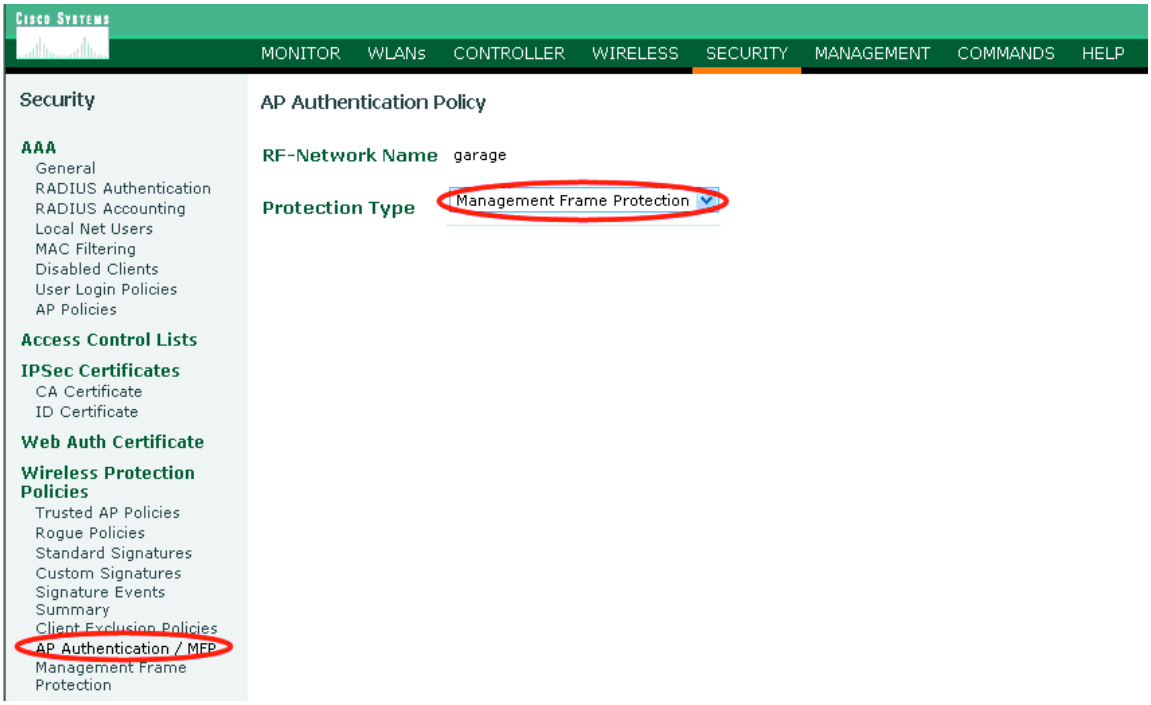

Figure 4-14 shows the enabling of MFP on the WLAN.

Figure 4-14 Enabling MFP per WLAN

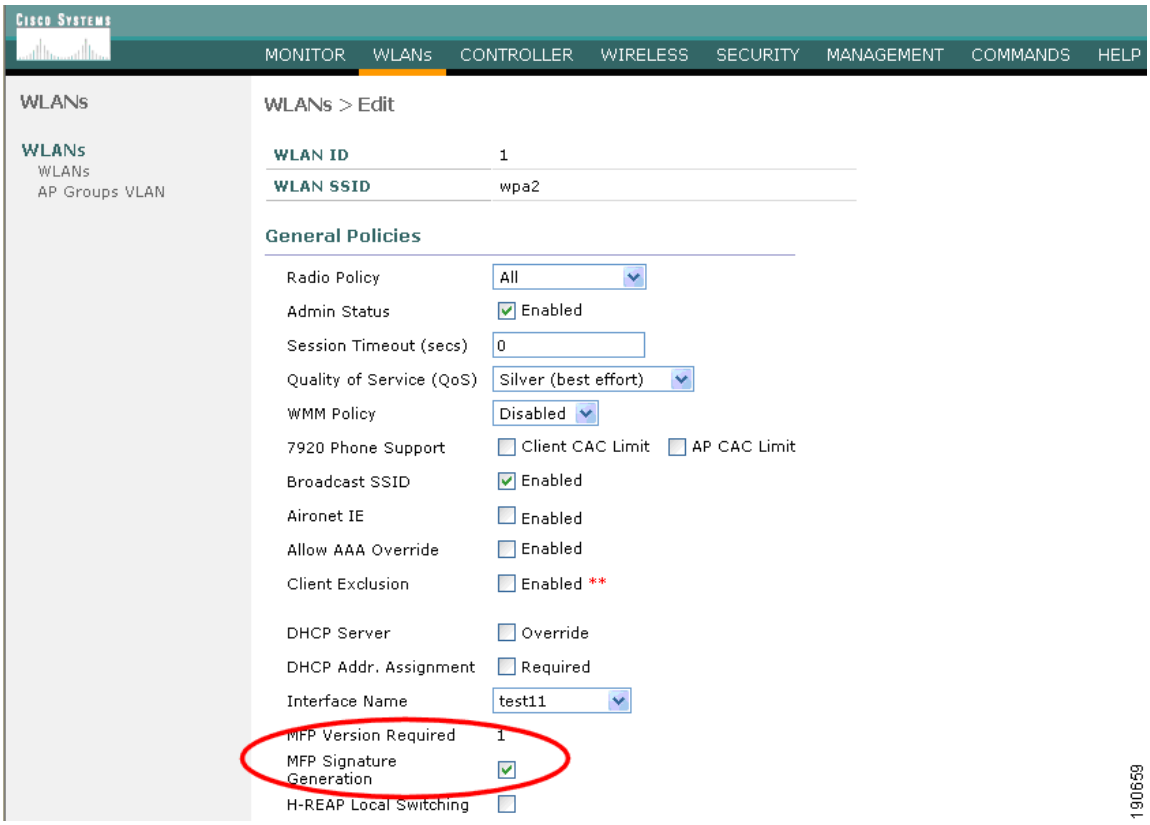

WLAN IDS

The WLC performs WLAN IDS analysis on all its connected WLANs APs, and reports detected attacks at the WLC as well to the WCS. This analysis is separate from the analysis that can be performed by a standalone network IDS system; it analyses 802.11 and WLC specific information that is not otherwise available to a network IDS.

The signature files used on the WLC are included in software releases, but can be updated independently through a signature file; these updated signatures are displayed in the Custom Signatures page.

Figure 4-15 shows the Standards Signatures page on the WLC.

Figure 4-15 Standard WLAN IDS Signatures

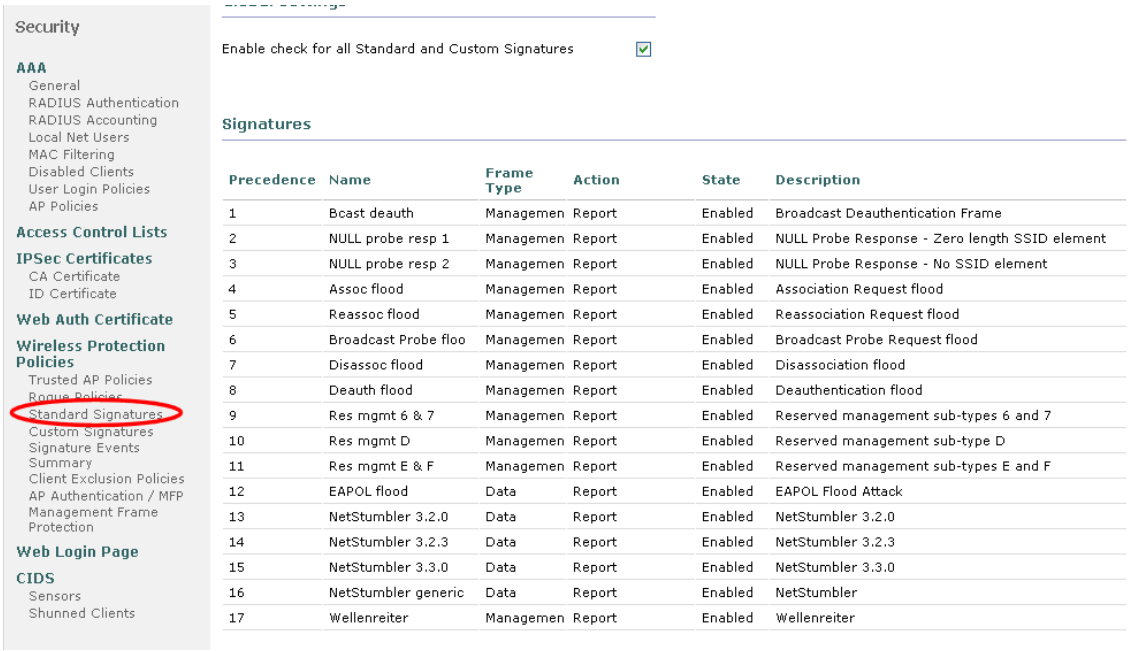

Client Security

The IDS features on the WLC provides alarm notifications for possible attacks on the WLAN network. Many of these are initiated by WLAN devices that are connected or attempting to connect to the WLAN network, and these cannot be blocked, only alarmed.

A separate set of client behaviors can be blocked, in addition to some behaviors that might warrant the disconnection of a client from WLAN network altogether. The blocking of clients is controlled through the client exclusion policy. Client exclusion is controlled on a per WLAN basis, as shown in Figure 4-16.

Figure 4-16 Enabling Client Exclusion

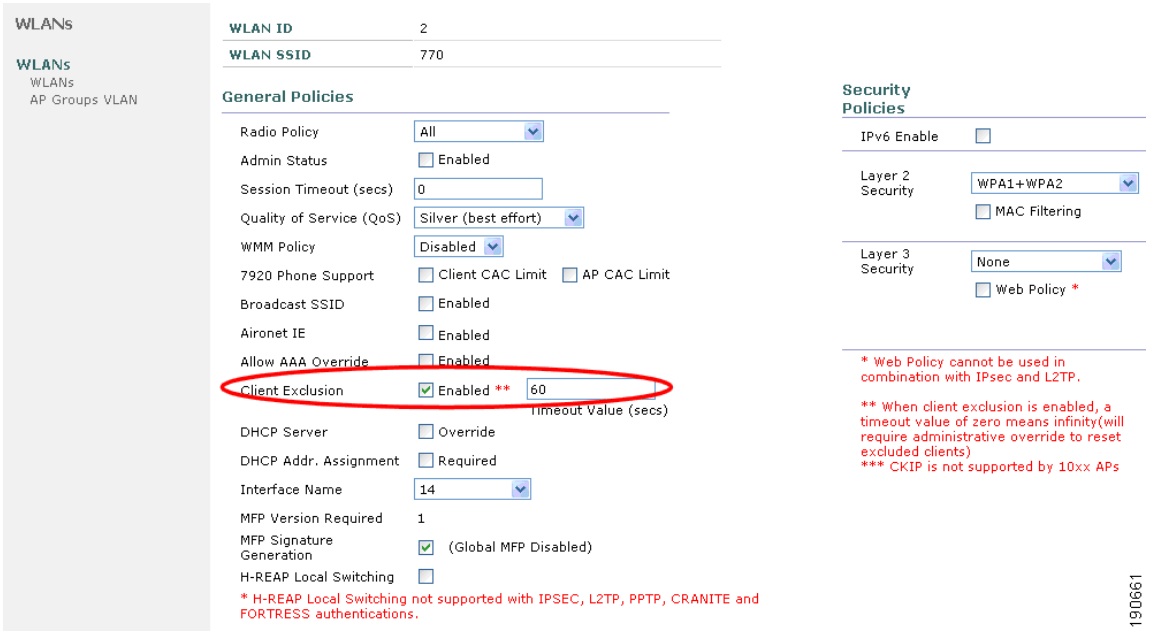

The suspect behaviors that cause client exclusion are configured on a per-controller basis, as shown in Figure 4-17.

Figure 4-17 Client Exclusion Policies

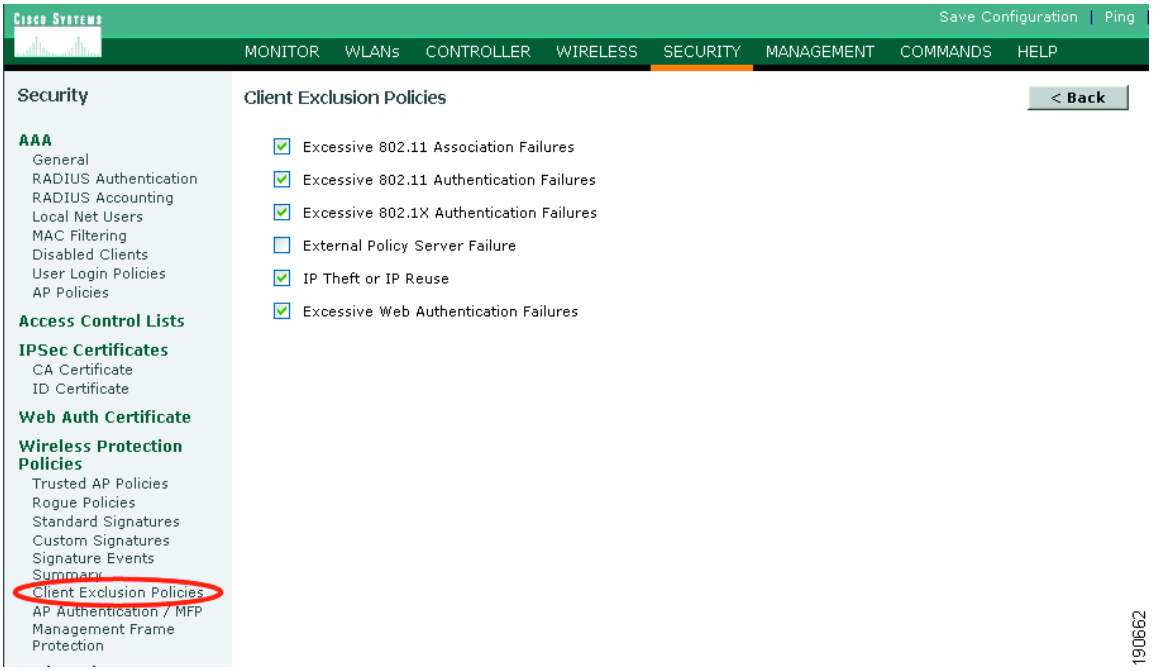

WLC Configuration

The three primary methods for configuring the WLC are HTTP, CLI, and SNMP. Each of these has security options. SNMP is covered later in the WCS chapter, but the primary means of securing the user interface is through the PKI encryption of HTTPS, and SSH. Figure 4-18 and Figure 4-19 show the configuration options for HTTP access and CLI access to the WLC. In each case, the encrypted or unencrypted communication mechanism can be selected.

Figure 4-18 HTTP Access to the WLC

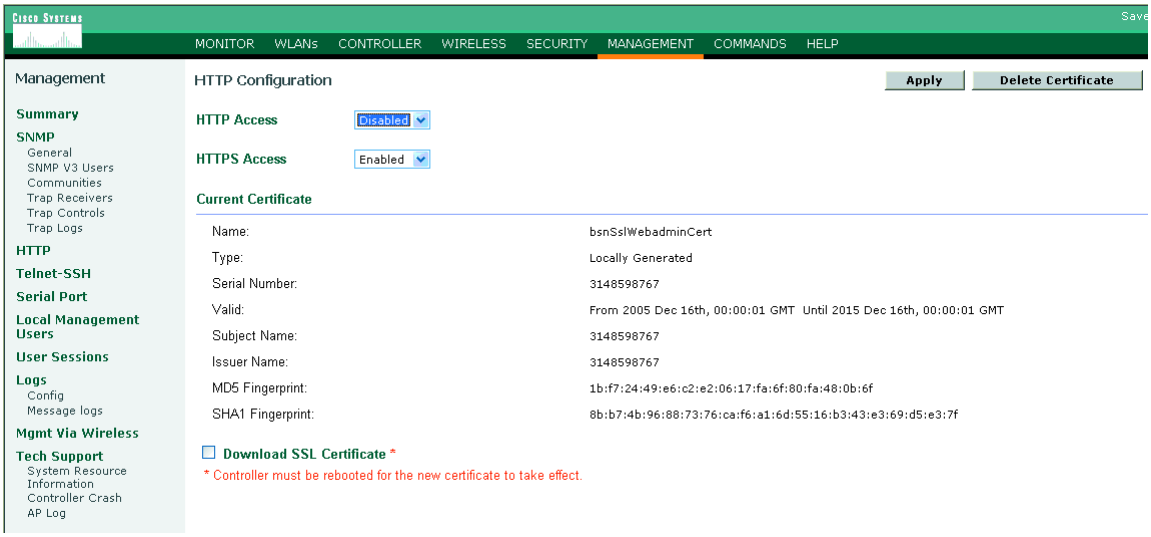

Figure 4-19 CLI Access to the WLC

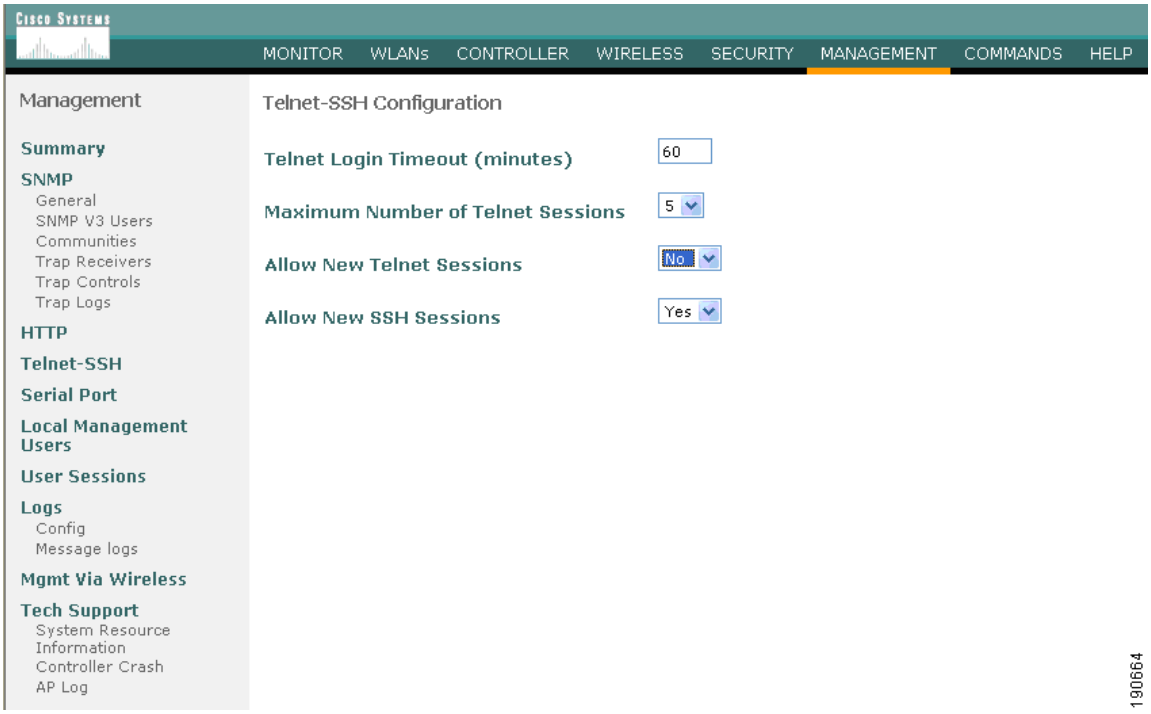

Management user authentication can be accomplished either through a local database or through a RADIUS server, as shown in Figure 4-20 and Figure 4-21.

Figure 4-20 Local Management Users

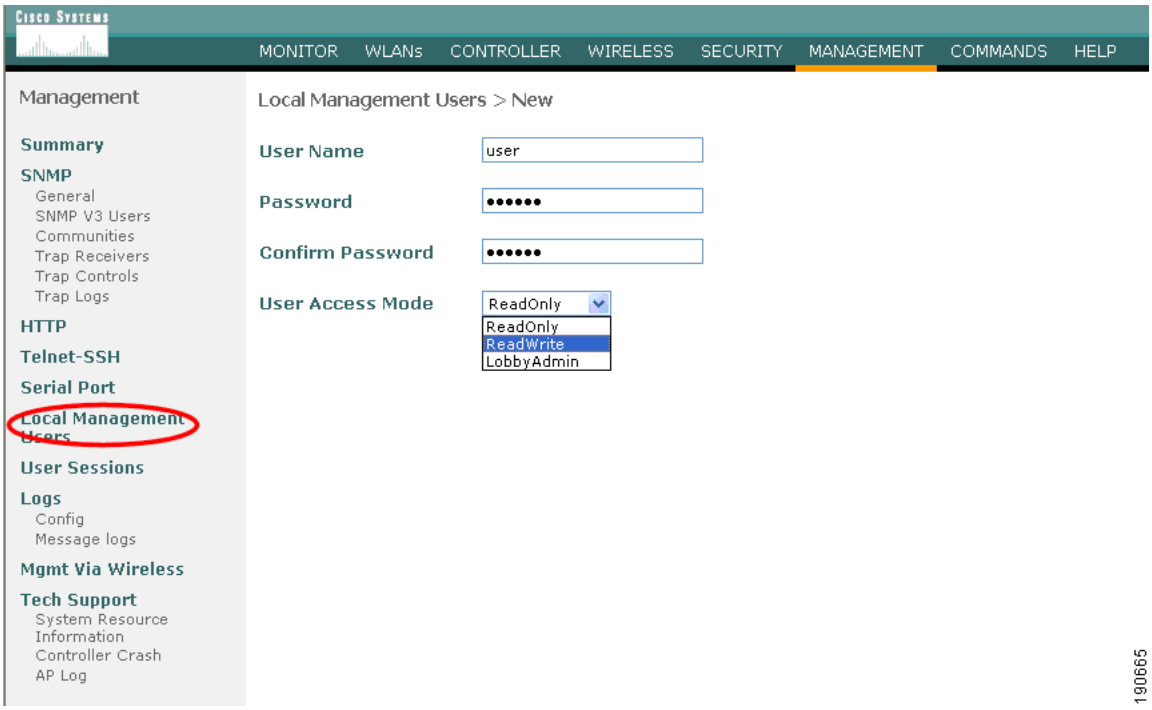

Figure 4-21 Management Users through RADIUS

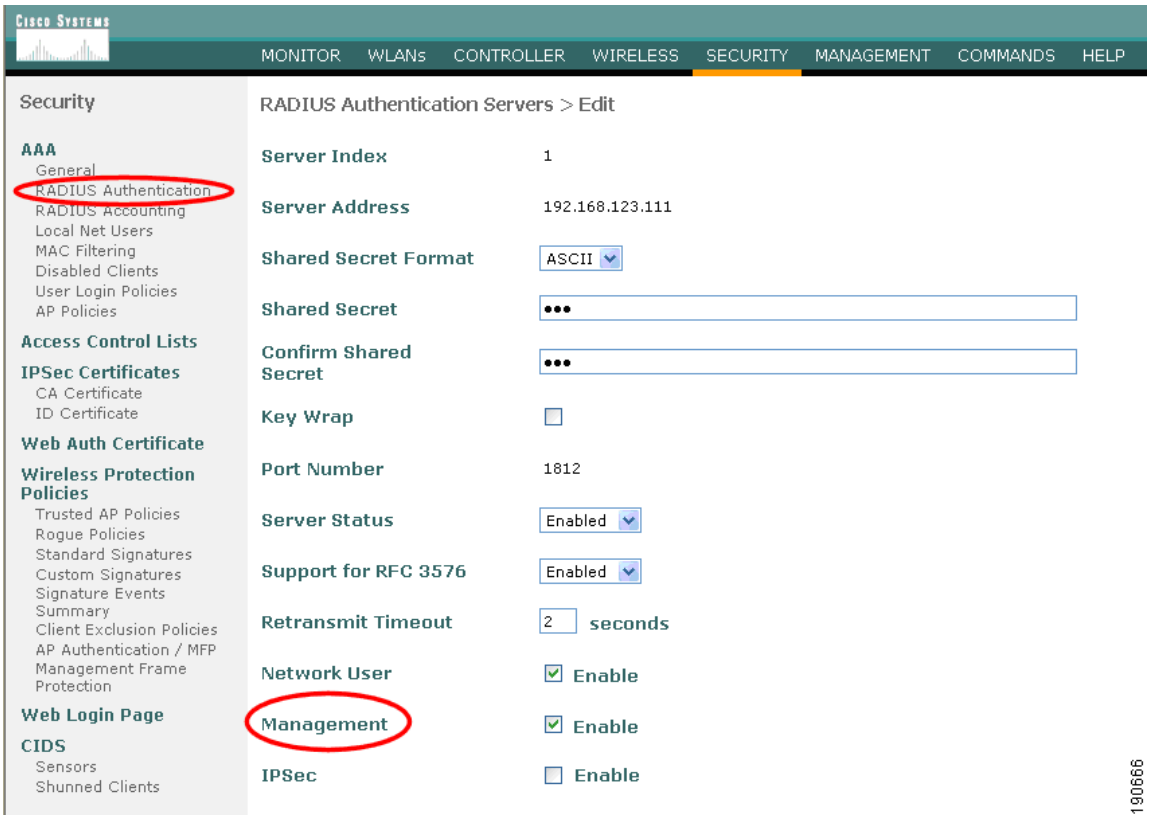

WLAN LAN Extension

The goal of a WLAN LAN extension network is for the WLAN access network to transparently provide the same applications and services as the wired access network. Each WLAN extension topic covered in this section addresses the following types of transparency:

- Security transparency—Do the selected security capabilities provide seamless WLAN network security equivalent to wired networks?
- Application transparency—Are the supported WLAN network applications identical to applications on a wired network?
- Performance transparency—Does the WLAN deliver application performance that matches wired network performance?
- User transparency—Are users of the WLAN forced to perform network-specific operations to use the WLAN?

WLAN LAN Extension 802.1x/EAP

This section presents WLAN Extension 802.1x/EAP deployment in terms of the following key topics:

- Security transparency
- Application transparency

- Performance transparency
- User transparency

An 802.1x/EAP implementation of WLAN LAN Extension operates at the link layer (Layer 2) to provide authentication, authorization, accounting, and encryption. Figure 4-22 shows the 802.1x/EAP WLAN.

The security level provided is beyond that provided on most wired networks, providing link layer encryption and Authentication, Authorization, and Accounting (AAA) access control. This is provided as follows:

- Authentication occurs between the client and the authentication server. Several EAP types (LEAP, EAP-FAST, EAP-TLS, PEAP) are supported, allowing the enterprise to choose the authentication type that best suits its needs.
- Encryption is at the link layer between the WLAN client and the AP. The encryption keys are automatically derived during the authentication process. Note that the LWAPP messages between the LWAPP AP and the controller are encrypted; but the client data, although LWAPP encapsulated is not encrypted.
- Authorization is controlled by the VLAN or interface membership given to the wireless client in combination with the access controls applied at the access router or switch terminating the VLAN or interface.
- Accounting is provided by the RADIUS accounting communicated by the WLC to the RADIUS server.

Figure 4-22 WLAN LAN Extension 802.1X/EAP

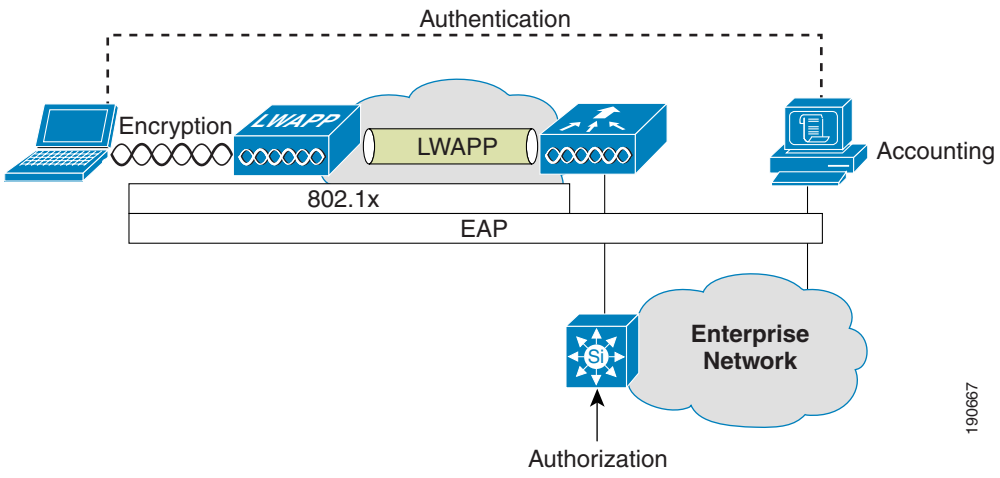

Application Transparency

The Cisco Unified Wireless architecture creates a virtual access/distribution network through the LWAPP protocol that aggregates WLAN traffic at the WLC. After the WLAN client traffic leaves the WLC, it is the same as wired traffic: subject to the same access control, queuing, and routing. This achieves the WLAN LAN extension goal of supporting the same applications as the wired network. Any inability to run applications from the wired network over the WLAN network would be the result of policies or the fundamental limitations of the WLAN, and not because of the 802.1x/EAP architecture. Figure 4-22 shows the Cisco Unified Wireless operation.

Performance Transparency

A WLAN has a lower bit rate and a lower throughput than most enterprise wired LANs. Therefore, providing equivalent performance for all applications over the WLAN can be a challenge. The strategy to minimize differences in application performance between the wired and WLAN network is to use the QoS tools available on the WLAN and the APs. Those applications identified as being sensitive to network throughput and delay can be classified and scheduled as required. Load balancing and admission control tools on the WLAN can optimize the usage of the available WLAN resources. After the user or device has been authenticated, there is an opportunity to apply identity based on QoS features.

User Transparency

The various EAP types in 802.1x/EAP allow enterprises to choose an authentication mechanism that best matches security requirements. This allows the integration of the 802.1x/EAP into existing user behavior. Many organizations enforce stronger authentication mechanisms on their WLAN networks (compared to wired networks), because of reduced physical security in the WLAN. Stronger authentication enforcement on wired networks is expected to catch up with WLAN networks, with organizations using 802.1x/EAP mechanisms to enhance wired network security.

WLAN LAN Extension IPsec

The use of IPsec VPN tunnels is an alternative to an 802.1x/EAP implementation. Network designers might choose this implementation over an 802.1x/EAP solution because of security policy reasons. IPsec is a well-established standard that is endorsed by a number of security organizations. IPsec is a regulatory requirement in some industries.

The primary advantage of an IPsec-based VPN solution is the encryption mechanism. IPsec includes support of Triple Data Encryption Standard (3DES) and AES encryptions, and wide deployment experience.

A WLAN LAN extension that makes use of IPsec is generally considered more difficult to implement than an 802.1x/EAP based solution, but the Cisco Unified WLAN Architecture greatly simplifies this deployment style by allowing untrusted WLAN VPN client traffic to be sent to a centralized location through LWAPP, tunnels to a WLC, or aggregated to multiple WLCs to an anchor WLC through the mobility anchor feature.

The network topology up to the VPN concentrator is considered untrusted, and an appropriate security policy must be created, configured, and maintained at all points that touch this untrusted network; for example, on the WLC, and the routers and switches between the WLC and the VPN concentrators. Some WLCs support VPN termination (440X model WLCs), and all WLCs can support VPN pass through, which is a mechanism to permit only VPN traffic destined for a external VPN concentrator on a certain WLAN.

Security Transparency

A WLAN LAN extension that uses IPsec provides AAA-equivalent features to that of 802.1x/EAP-based solutions (see Figure 4-23). Key elements are as follows:

- Authentication occurs between the client and the VPN concentrator. Multiple authentication types are supported within the IPsec framework.

- Encryption is at the network layer using 3DES or AES, and is negotiated between the client and the VPN concentrator. In addition to the inherent WLAN LAN extension IPsec security features associated with this implementation, VPN capabilities provide additional AAA-related security capabilities.
- Authorization is controlled by the VPN concentrator and is determined at the time of authentication. Policy is provided by the authentication server.
- Accounting is provided by RADIUS accounting software on both the VPN concentrator and the authentication server.

Figure 4-23 WLAN LAN Extension IPsec VPN

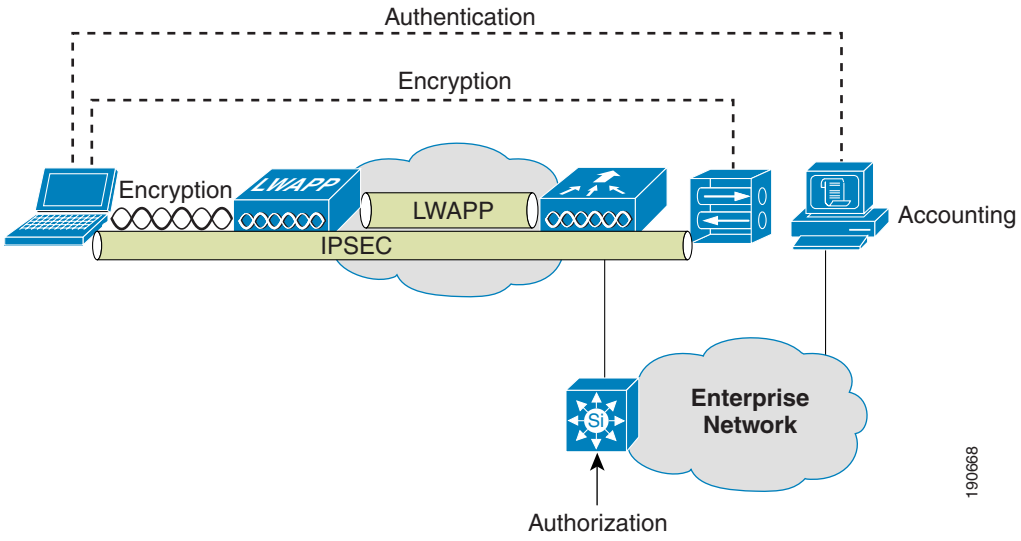

Application Transparency

As can be seen in Figure 4-23, WLAN traffic is transported over an IPsec tunnel to the VPN concentrator.

This can affect application transparency:

- Protocol limitations—Only the IP protocol is supported; the network is not multi-protocol
- Address translation—The IPsec client performs a form of address translation between its local IP address and that allocated by the VPN concentrator. This can impact the operation of some applications.
- No multicast—The connection to the VPN concentrator is point-to-point. Multicast applications are not supported.

Performance Transparency

Providing equivalent performance for all applications over the WLAN can be a challenge, because a WLAN has a lower bit rate and a lower throughput than most enterprise wired LANs. The use of IPsec VPN tunnels introduces some additional considerations:

- MTU size—The MTU size of packets must be adjusted to incorporate IPsec overhead.

- Processing overhead—Clients incur processing overhead from IPsec VPN. However, this should not be noticeable on most target platforms.
- Traffic classification and QoS considerations—Type of service (ToS) and differentiated services code point (DSCP) values are projected from client packets into the IPsec packets. As a result, QoS preference can be acted on, but no classification of traffic is possible while the traffic is IPsec encrypted.
- Traffic scheduling—All queuing at the VPN concentrator is handled on a first-in-first-out basis.

User Transparency

The Cisco IPsec VPN client has a number of features that aid user transparency, thereby providing an equivalent user experience when compared to that of 802.1x/EAP solutions:

- Auto Initiation—The VPN client can be configured to automatically launch for particular address ranges. In an enterprise, this would be configured to launch within the enterprise WLAN address ranges.
- OS Integration—The VPN client can capture user name and password information at login and use these as part of the VPN client login. This is similar to the process used in EAP-Cisco. As an alternative, the VPN client can use stored certificates associated with a specific user, similar to EAP-TLS. These features coupled with Auto Initiation should provide a high level of user transparency.

WLAN Static Keys

Static key implementations are not recommended for general purpose WLAN LAN extension networks because of known weaknesses in the WEP encryption algorithms, and because of the difficulty in the configuration and maintenance of static keys for WEP or other stronger encryption schemes. Certain client devices are capable of supporting static WEP keys only (see Figure 4-24). These clients should be put on a separate WLAN VLAN or interface and have their authorization limited to addresses and protocols specific to the application supported by the Static WEP client. If possible, WPA-PSK or WPA2-PSK should be used in place of WEP because these mechanisms address the known weaknesses in the WEP encryption system

Figure 4-24 WLAN Static WEP Keys

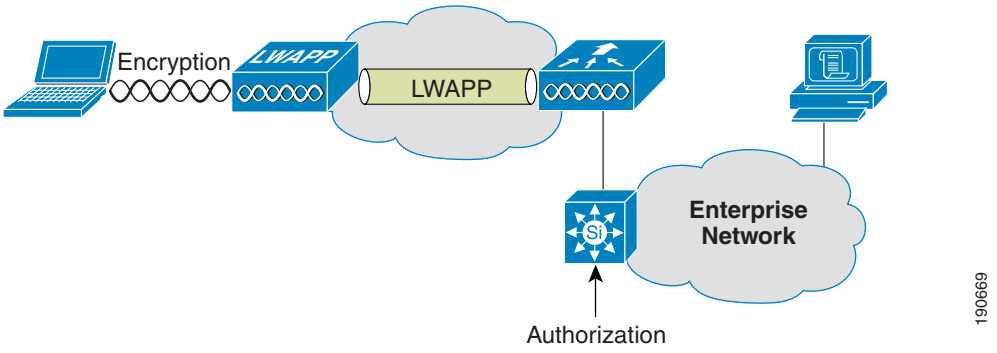

Security Transparency

Some security issues related to static key implementations are as follows:

- Weak authentication—Any hardware device with a matching configuration and key can join the network. The Static key authenticates a group of devices, never individual users. MAC filtering can be added, but MAC addresses are sent in the clear, and can be spoofed.
- Encryption limitation—Encryption is at the link layer between the WLAN client and the AP. The current encryption mechanisms available are WEP, WPA-PSK, or WPA2-PSK. If possible, WPA-PSK or WPA2-PSK should be used.
- Authorization limitation—Authorization is controlled by the VLAN membership associated with the SSID, or assigned through MAC filtering.
- Accounting is not available.

Application Transparency

As illustrated in Figure 4-24, the WLAN connects at the access/distribution layer. When the WLAN client traffic leaves the WLC, it is the same as wired network traffic and subject to the same access control, queuing, and routing. WLAN Static key solutions should be limited to the specialized applications that the Static WEP client supports. The network would appear transparent to this application, but to all other applications access should be blocked.

Performance Transparency

To minimize differences in application performance between the wired and WLAN network, use the QoS tools available on the WLAN, the APs, and WLC. Those applications identified as being sensitive to network throughput and delay can be classified and scheduled as required. Load balancing and admission control tools on the WLAN can optimize the usage of the available WLAN resources. Because Static WEP performs no user authentication, no user-based QoS policies can be applied, but MAC-based QoS policies are possible.

User Transparency

Static WEP requires no authentication and should be transparent to the supported applications and users. The static WEP key becomes an issue only for the user if required to change it.

Cisco Unified WLAN Architecture Considerations

The Cisco Unified WLAN architecture has features that can enhance solution transparency. The following section details some of the specific considerations. For more information, see the following URL:

http://www.cisco.com/en/US/products/ps6366/prod_technical_reference09186a00806cfa96.html

Security Transparency

The features offered by Cisco Unified WLAN architecture do not directly impact security transparency because the architecture supports all the existing security models. An integrated WLC solution, such as WISM, can make it easier to implement various security solutions through integration with IOS features on that platform.

Application Transparency

PKC and CCKM enable WLAN clients to quickly roam between APs. The WLC caches session credentials (security keys) derived for a client session and uses them for re-authentication and re-keying when a client roams, within the mobility group. Caching this information rather than forcing the client to do a full authentication reduces the authentication time and therefore the total time required for roaming. This can enhance application transparency because the impact of roaming is reduced and less likely to impact either the application or the user.

Performance Transparency

The Cisco Unified WLAN Architecture has been designed to use and maintain the QoS features used in neighboring wired networking platforms.

User Transparency

Cisco Unified WLAN Architecture is compatible with all other WLAN client solutions, and therefore does not have an adverse impact on user transparency.

Note Cisco Unified WLAN architecture is compatible with CCXv4 with the 4.0 controller software release.

EAP Considerations for High Availability ACS Architecture

As a centralized authentication server, Cisco Secure ACS introduces RADIUS-based AAA capabilities to an enterprise network for both wired and WLAN networks. Implementing ACS redundancy and reliability is meant to address two issues:

- The ACS server should not represent a single point of failure.
- A network failure should not impact a user's ability to log on.

The first issue is a good reason to replicate the ACS database to a secondary server, allowing for failover and maintenance. This redundancy configuration should be implemented in almost all cases. The second issue is an instance in which it is critical to use the local WLAN even in the event of a network failure preventing access to a remote ACS server. Implementation of this second use of replication depends on the application architecture of the enterprise. For example, if the applications that the users want to reach are also remote, little is to be gained by being able to use the WLAN.

One issue that should also be considered in RADIUS planning is the impact that WAN latency can have on authentication. This is especially true in (non key caching) re-authentication scenarios where a full authentication back to the RADIUS server is required when a client roams between APs and thereby adds

latency to the client roam. In cases where clients roaming times need to be minimized, key-caching mechanisms such as PKC or CCKM should be considered. These mechanisms have the advantage of requiring full RADIUS authentication only initially and using the cached key when a client roams, reducing the client roam times, and reducing the load on the WAN and RADIUS server.

ACS Architecture

The ACS deployment strategy must consider how the entire enterprise identity system will be structured, rather than just the campus. A key consideration is the location of directory databases. It is essential that the ACS strategy reflect an approach in which the elements of the ACS architecture are carefully analyzed, designed, and implemented for authentication systems associated with the directory architecture of the organization. The assessment of the directory architecture is the starting point for the ACS deployment strategy. In an ideal situation, the existing infrastructure can provide the user names, passwords, and profiles to the ACS servers. That is, the ACS acts as a AAA interface between users and the directory system, and placement of ACS servers needs to align with the placement of directory resources.

In deploying multiple ACS servers, the ACS replication service can be used to replicate a master ACS with slave ACSs in the network (the replication is master/slave).

Sample Architecture

Figure 4-25 shows an example of what ACS architecture might look like. Campus A holds the authoritative ACS database server. This server is replicated to the other enterprise ACS servers. WLCs communicate to the two local ACS servers.

Campus B, because of its size and distance from Campus A, has opted for another two ACS servers, thus providing its own backup. Campus C, being smaller and closer to Campus A, has opted to have only one server, and relies on Campus A for backup. The branch offices use the ACS servers that are the shortest network distance from them.

Figure 4-25 Sample ACS Architecture

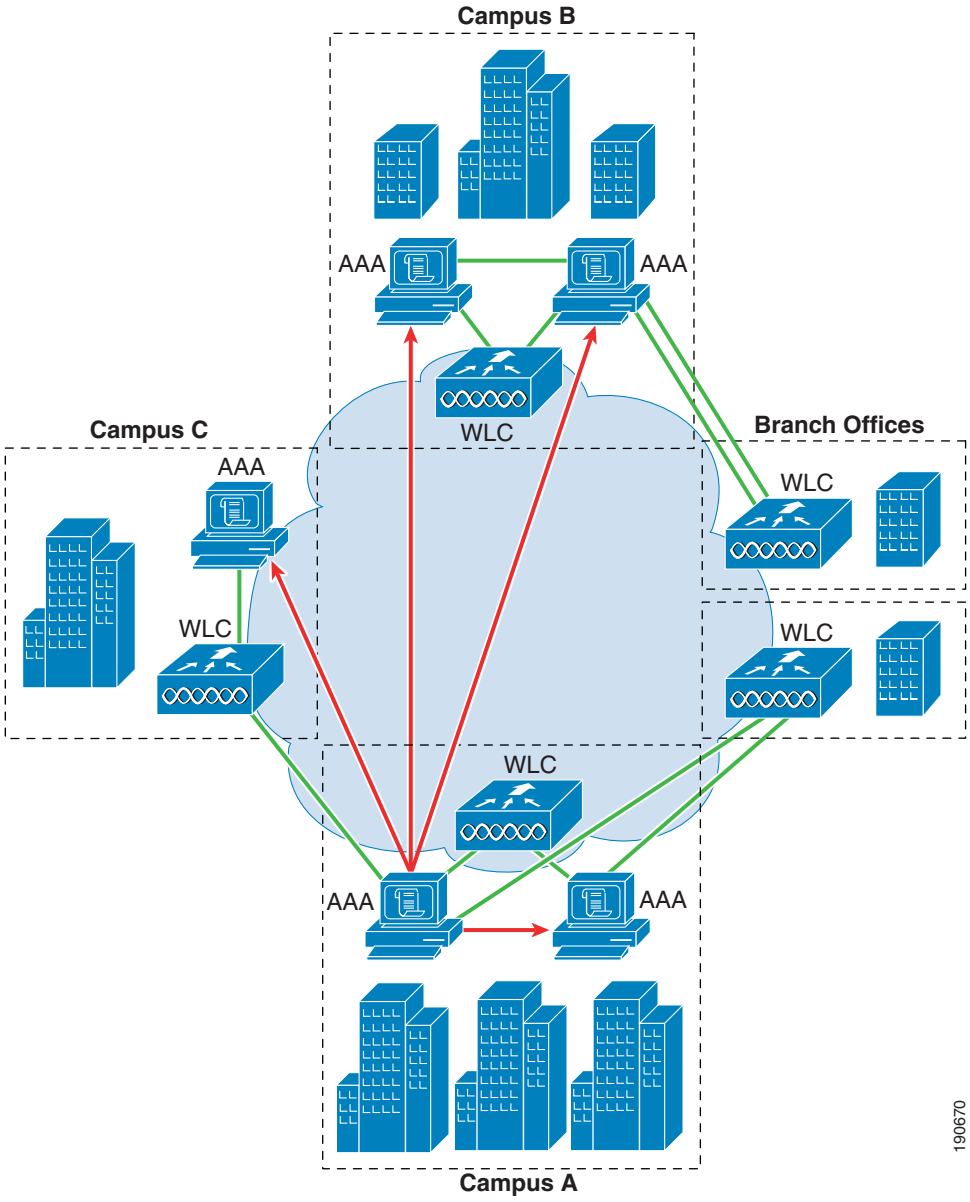

CHAPTER 5

Cisco Unified Wireless QoS

Introduction

This chapter describes quality of service (QoS) in the context of WLAN implementations. This chapter describes WLAN QoS in general, but does not provide in-depth coverage on topics such as security, segmentation, and voice over WLAN (VoWLAN), although these topics have a QoS component. This chapter also provides information on the features of the Cisco Centralized WLAN Architecture.

This chapter is intended for those who are tasked with designing and implementing enterprise WLAN deployments using the Cisco Unified Wireless technology.

QoS Overview

QoS refers to the capability of a network to provide better service to selected network traffic over various network technologies. QoS technologies provide the following benefits:

- Provides building blocks for business multimedia and voice applications used in campus, WAN, and service provider networks
- Allows network managers to establish service level agreements (SLAs) with network users
- Enables network resources to be shared more efficiently and expedites the handling of mission-critical applications
- Manages time-sensitive multimedia and voice application traffic to ensure that this traffic receives higher priority, greater bandwidth, and less delay than best-effort data traffic

With QoS, bandwidth can be managed more efficiently across LANs, including WLANs and WANs. QoS provides enhanced and reliable network service by:

- Supporting dedicated bandwidth for critical users and applications
- Controlling jitter and latency (required by real-time traffic)
- Managing and minimizing network congestion
- Shaping network traffic to smooth the traffic flow
- Setting network traffic priorities

Wireless QoS Deployment Schemes

In the past, WLANs were mainly used to transport low-bandwidth, data-application traffic. Currently, with the expansion of WLANs into vertical (such as retail, finance, and education) and enterprise environments, WLANs are used to transport high-bandwidth data applications, in conjunction with time-sensitive, multimedia applications. This requirement led to the necessity for wireless QoS.

Several vendors, including Cisco, support proprietary wireless QoS schemes for voice applications. To speed up the rate of QoS adoption and to support multi-vendor time-sensitive applications, a unified approach to wireless QoS is necessary. The IEEE 802.11e working group within the IEEE 802.11 standards committee has completed the standard definition, but adoption of the 802.11e standard is in its early stages, and as with many standards there are many optional components. Just as occurred with 802.11 security in 802.11i, industry groups such as the WiFi Alliance and industry leaders such as Cisco are defining the key requirements in WLAN QoS through their WMM and CCX programs, ensuring the delivery of key features and interoperation through their certification programs.

Cisco Unified Wireless Products support Wi-Fi MultiMedia (WMM), a QoS system based on the IEEE 802.11e draft that has been published by the Wi-Fi Alliance. An example deployment of wireless QoS based on Cisco Unified Wireless technology features is shown in Figure 5-1.

Figure 5-1 Wireless QoS Deployment Example

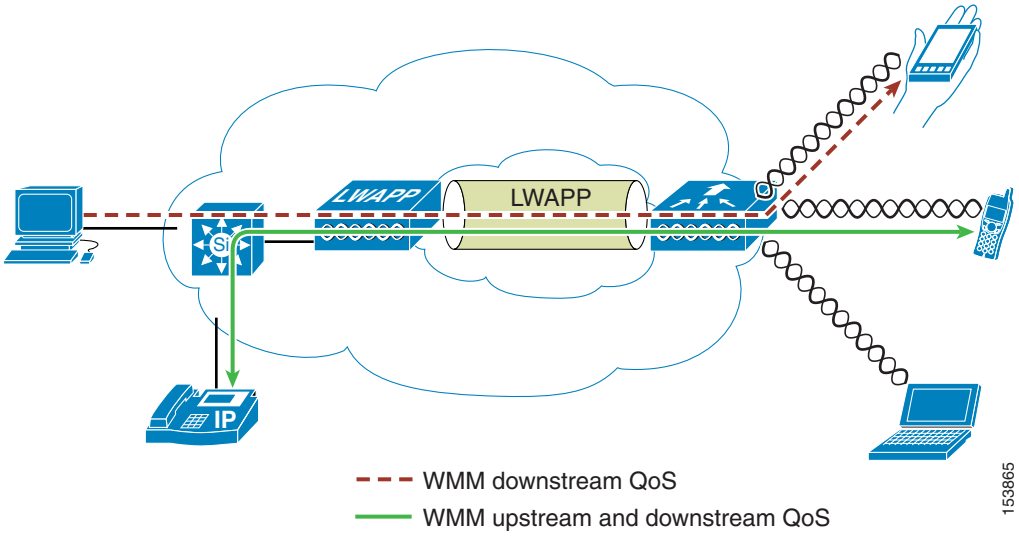

QoS Parameters

QoS is defined as the measure of performance for a transmission system that reflects its transmission quality and service availability. Service availability is a crucial element of QoS. Before QoS can be successfully implemented, the network infrastructure must be highly available. The network transmission quality is determined by latency, jitter, and loss, as shown in Table 5-1.

Table 5-1 QoS Parameters

Transmission Quality	Description
Latency	Latency (or delay) is the amount of time it takes for a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time period is called the *end-to-end delay* and can be divided into two areas: fixed network delay and variable network delay. *Fixed network delay* includes encoding and decoding time (for voice and video), and the finite amount of time required for the electrical or optical pulses to traverse the media en route to their destination. *Variable network delay* generally refers to network conditions, such as congestion, that can affect the overall time required for transit.
Jitter	Jitter (or delay-variance) is the difference in the end-to-end latency between packets. For example, if one packet requires 100 mSec to traverse the network from the source endpoint to the destination endpoint and the next packet requires 125 mSec to make the same trip, then the jitter is calculated as 25 mSec.
Loss	Loss (or packet loss) is a comparative measure of packets successfully transmitted and received to the total number that were transmitted. Loss is expressed as the percentage of packets that were dropped.

Upstream and Downstream QoS

Figure 5-2 illustrates the definition of QoS radio *upstream* and *downstream*.

The notations in Figure 5-2 refer to the following:

- *Radio downstream* QoS refers to the traffic leaving the AP and traveling to the WLAN clients. Radio downstream QoS is the primary focus of this chapter, because this is still the most common deployment. The client upstream QoS depends on the client implementation.
- *Radio upstream* QoS refers to traffic leaving the WLAN clients and traveling to the AP. WMM provides upstream QoS for WLAN clients supporting WMM.
- *Network downstream* refers to traffic leaving the WLC traveling to the AP. QoS can be applied at this point to prioritize and rate limit traffic to the AP. Configuration of Ethernet downstream QoS is not covered in this chapter.
- *Network upstream* refers to traffic leaving the AP, traveling to the WLC. The AP classifies traffic from the AP to the upstream network according to the traffic classification rules of the AP.

Figure 5-2 Upstream and Downstream QoS

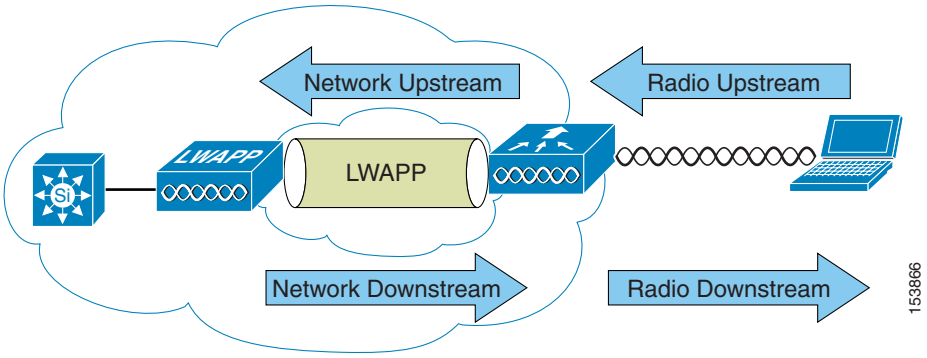

QoS and Network Performance

The application of QoS features might not be easily detected on a lightly loaded network. If latency, jitter, and loss are noticeable when the media is lightly loaded, it indicates a system fault, poor network design, or that the latency, jitter, and loss requirements of the application are not a good match for the network.

QoS features start to impact application performance as the load on the network increases. QoS works to keep latency, jitter, and loss for selected traffic types within acceptable boundaries.

When providing only radio downstream QoS from the AP, radio upstream client traffic is treated as best-effort. A client must compete with other clients for upstream transmission as well as competing with best-effort transmission from the AP. Under certain load conditions, a client can experience upstream congestion and the performance of QoS-sensitive applications might be unacceptable despite the QoS features on the AP.

Ideally upstream and downstream QoS can be operated either by using WMM on both the AP and WLAN client, or by using WMM and a client's proprietary implementation.

Note Even without WMM support on the WLAN client, the Cisco Unified Wireless solution is able to provide network prioritization in both network upstream and network downstream situations.

Note WLAN client support for WMM does not mean that the client traffic automatically benefits from WMM. The applications looking for the benefits of WMM assign an appropriate priority classification to their traffic, and the operating system needs to pass that classification to the WLAN interface. In purpose-built devices, such as VoWLAN handsets, this is done as part of the design, but if implementing on a general purpose platform, such as a PC, application traffic classification and OS support must be implemented before the WMM features can be used to good effect.

802.11 DCF

Data frames in 802.11 are sent using the Distributed Coordination Function (DCF), which is composed of two main components:

- Interframe spaces (SIFS, PIFS, and DIFS)

- Random backoff (contention window)

DCF is used in 802.11 networks to manage access to the RF medium. A baseline understanding of DCF is necessary to deploy 802.11e-based enhanced distributed channel access (EDCA). See the IEEE 802.11 specification for more information on DCF at the following URL: http://ieeexplore.ieee.org/xpl/standardstoc.jsp?isnumber=14251&isYear=1997

Interframe Spaces

802.11 currently defines three interframe spaces, as shown in Figure 5-3:

- Short interframe space (SIFS) 10 µs
- Point interframe space (PIFS) SIFS + 1 x slot time = 30 µs
- Distributed interframe space (DIFS) 50 µs SIFS + 2 x slot time = 50 µs

The interframe spaces, SIFS, PIFS, and DIFS allow 802.11 to control which traffic gets first access to the channel after carrier sense declares the channel to be free.

Figure 5-3 Interframe Spaces (IFS)

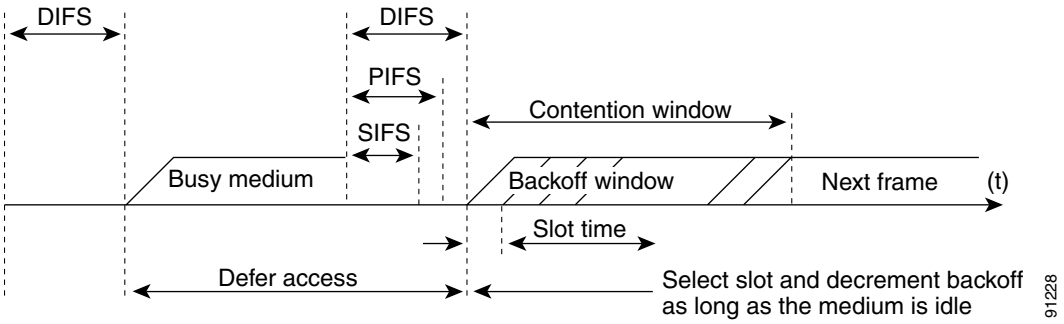

Random Backoff

When a data frame using Distributed Coordination Function (DCF), shown in Figure 5-4, is ready to be sent, it goes through the following steps:

1. Generates a random backoff number between 0 and a minimum Contention Window (CWmin).
2. Waits until the channel is free for a DIFS interval.
3. If the channel is still free, begins to decrement the random backoff number, for every slot time (20 µs) the channel remains free.
4. If the channel becomes busy, such as another station getting to 0 before your station, the decrement stops and steps 2 through 4 are repeated.
5. If the channel remains free until the random backoff number reaches 0, the frame can be sent.

Figure 5-4 shows a simplified example of how the DCF process works. In this simplified DCF process, no acknowledgements are shown and no fragmentation occurs.

Figure 5-4 Distributed Coordination Function Example

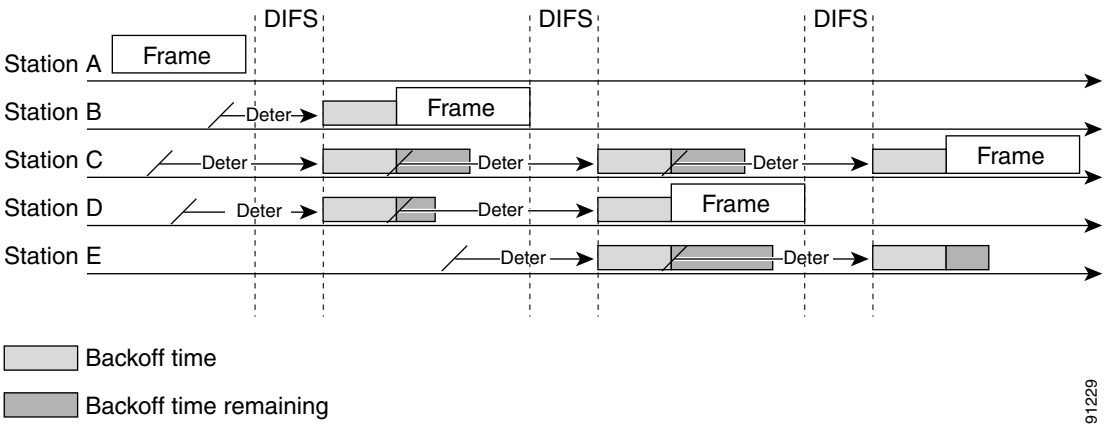

The DCF steps illustrated in Figure 5-4 are as follows:

1. Station A successfully sends a frame, and three other stations also want to send frames but must defer to Station A traffic.
2. After Station A completes the transmission, all the stations must still defer for the DIFS. When the DIFS is complete, stations waiting to send a frame can begin to decrement the backoff counter, once every slot time, and can send their frame.
3. The backoff counter of Station B reaches zero before Stations C and D, and therefore Station B begins transmitting its frame.
4. When Station C and D detect that Station B is transmitting, they must stop decrementing the backoff counters and defer until the frame is transmitted and a DIFS has passed.
5. During the time that Station B is transmitting a frame, Station E gets a frame to transmit, but because Station B is sending a frame, it must defer in the same manner as Stations C and D.
6. When Station B completes transmission and the DIFS has passed, stations with frames to send begin to decrement the backoff counters. In this case, Station D's backoff counter reaches zero first and it begins transmission of its frame.
7. The process continues as traffic arrives on different stations.

CWmin, CWmax, and Retries

DCF uses a contention window (CW) to control the size of the random backoff. The contention window is defined by two parameters:

- aCWmin
- aCWmax

The random number used in the random backoff is initially a number between 0 and aCWmin. If the initial random backoff expires without successfully sending the frame, the station or AP increments the retry counter, and doubles the value random backoff window size. This doubling in size continues until the size equals aCWmax. The retries continue until the maximum retries or time to live (TTL) is reached. This process of doubling the backoff window is often referred to as a *binary exponential backoff*, and is illustrated in Figure 5-5.

Figure 5-5 Growth in Random Backoff Range with Retries

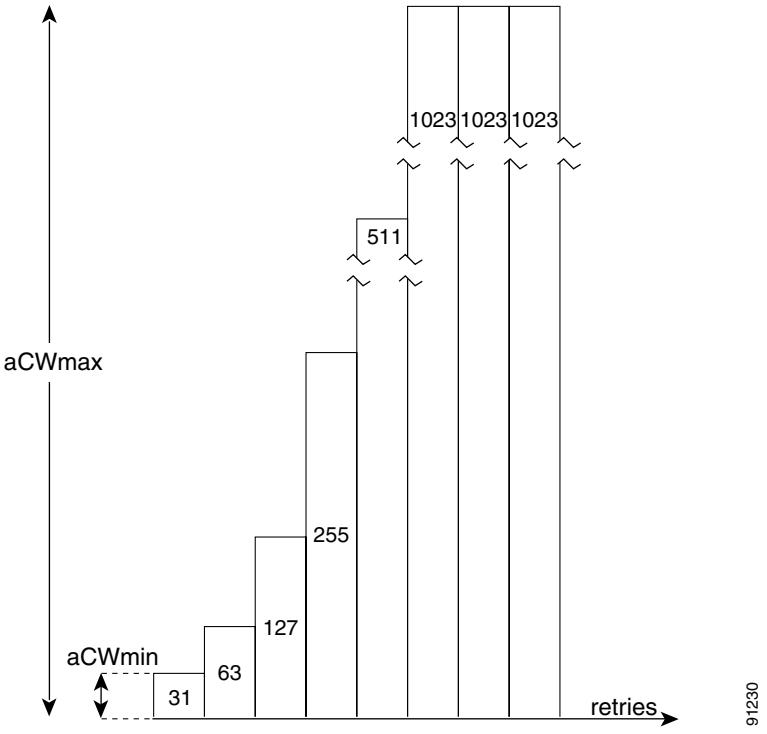

Wi-Fi Multimedia

This section describes three Wi-Fi Multimedia (WMM) implementations:

- WMM access
- WMM power save
- WMM access control

WMM Access

WMM is a Wi-Fi Alliance certification of support for a set of features from the 802.11e draft. This certification is for both clients and APs, and certifies the operation of WMM. WMM is primarily the implementation of the EDCA component of 802.11e. Additional Wi-Fi certifications are planned to address other components of the 802.11e.

Figure 5-6 shows the principle behind EDCF, where different interframe spacing and CWmin and CwMax values are applied per traffic classification. Different traffic types can wait different interface spaces before counting down their random backoff, and the CW value used to generate the random backoff number also depends on the traffic classification. High priority traffic has small interframe space and a small CWmin value, giving as short random backoff, whereas best-effort traffic has a longer interframe space and large CWmin value that on average gives a large random backoff number.

Figure 5-6 Access Category Timing

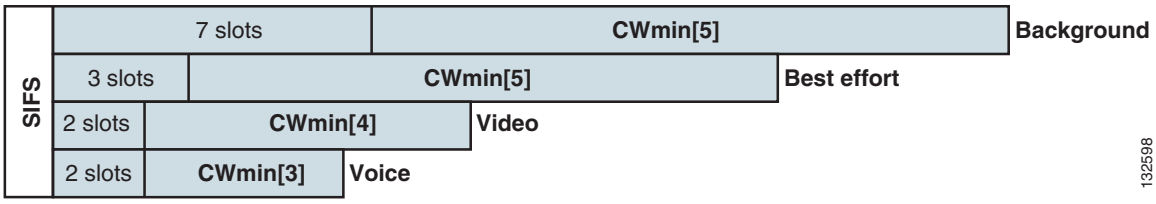

WMM Classification

WMM uses the 802.1p classification scheme developed by the IEEE (which is now a part of the 802.1D specification).

This classification scheme has eight priorities, which WMM maps to four access categories: AC_BK, AC_BE, AC_VI, and AC_VO. These access categories map to the four queues required by a WMM device, as shown in Table 5-2.

Table 5-2 802.1p and WMM Classification

Priority	802.1 Priority (=User Priority)	802.1p Designation	Access Category	WMM Designation
Lowest	1	BK Background	AC_BK	Background
	2	-Spare		
	0	BE Best-effort		
	3	EE Excellent Effort	AC_BE	Best-effort
	4	CL Control Load		
	5	VI Video <100ms	AC_VI	Video
	6	VO Voice <10ms	AC_VO	Voice
Highest	7	NC Network Control “must get there”		

Figure 5-7 shows the WMM data frame format. Note that even though WMM maps the eight 802.1p classifications to four access categories, the 802.11D classification is sent in the frame.

Note The WMM and IEEE 802.11e classifications are different from the classifications recommended and used in the Cisco network, which are based on IETF recommendations. The primary difference in classification is the demoting of voice and video traffic to 5 and 4, respectively. This allows the 6 classification to be used for Layer 3 network control. To be compliant with both standards, the Cisco Unified Wireless solution performs a conversion between the various classification standards when the traffic crosses the wireless-wired boundary.

Figure 5-7 WMM Frame Format

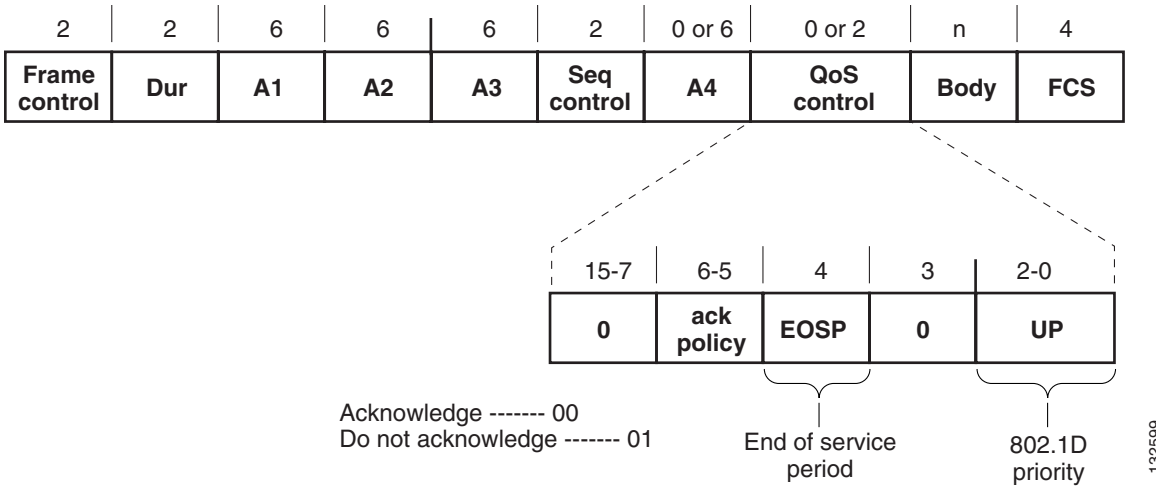

WMM Queues

Figure 5-8 shows the queuing performed on a WMM client or AP. There are four separate queues, one for each of the access categories. Each of these queues contends for the wireless channel in a similar manner to the DCF mechanism described previously, with each of the queues using different interframe space, CWmin, and CWmax values. If more than one frame from different access categories collide internally, the frame with the higher priority is sent, and the lower priority frame adjusts its backoff parameters as though it had collided with a frame external to the queuing mechanism. This system is called enhanced distributed channel access (EDCA).

Figure 5-8 WMM Queues

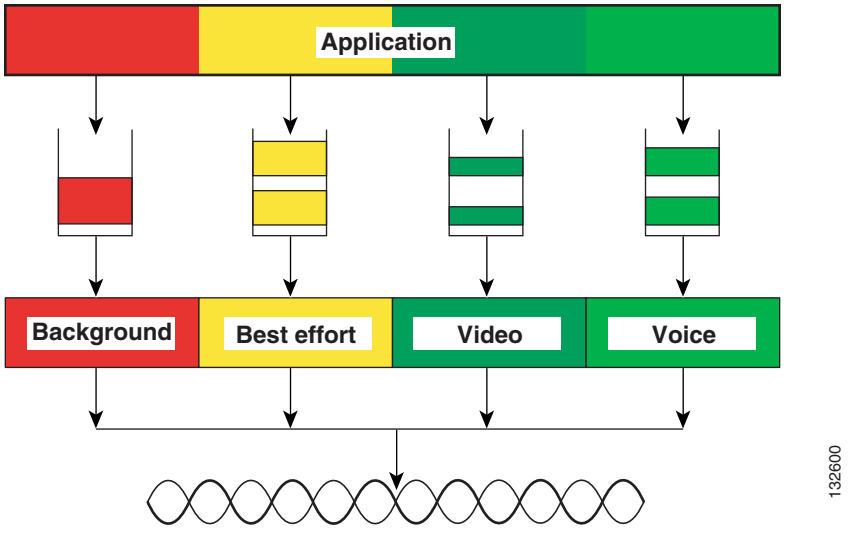

EDCA

The EDCA process is illustrated in Figure 5-9, using data from Figure 5-10, and follows this sequence:

1. While Station X is transmitting its frame, three other stations determine that they must send a frame. Each station defers because a frame was already being transmitted, and each station generates a random backoff.
2. Because station Voice has a traffic classification of voice, it has an arbitrated interframe space (AIFS) of 2, and uses an initial CWmin of 3, and therefore must differ the countdown of its random backoff for 2 slot times, and has a short random backoff value.
3. Best-effort has an AIFS of 3 and a longer random backoff time, because its CWmin value is 5.
4. Voice has the shortest random backoff time, and therefore starts transmitting first. When Voice starts transmitting, all other stations defer.
5. After Voice Station finishes transmitting, all stations wait their AIFS, then begin to decrement the random backoff counters again.
6. Best-effort then completes decrementing its random backoff counter and begins transmission. All other stations defer. This can happen even though there might be a voice station waiting to transmit. This shows that best-effort traffic is not starved by voice traffic because the random backoff decrementing process eventually brings the best-effort backoff down to similar sizes as high priority traffic, and that the random process might, on occasion, generate a small random backoff number for best-effort traffic.
7. The process continues as other traffic enters the system. The access category settings shown in Table 5-3 and Table 5-4 are, by default, the same for an 802.11a radio, and are based on formulas defined in WMM.

Note Table 5-3 refers to the parameter settings on a client, which are slightly different from the settings for an AP. This is because an AP is expected to have multiple clients, and therefore needs to send frames more often.

Figure 5-9 EDCA Example

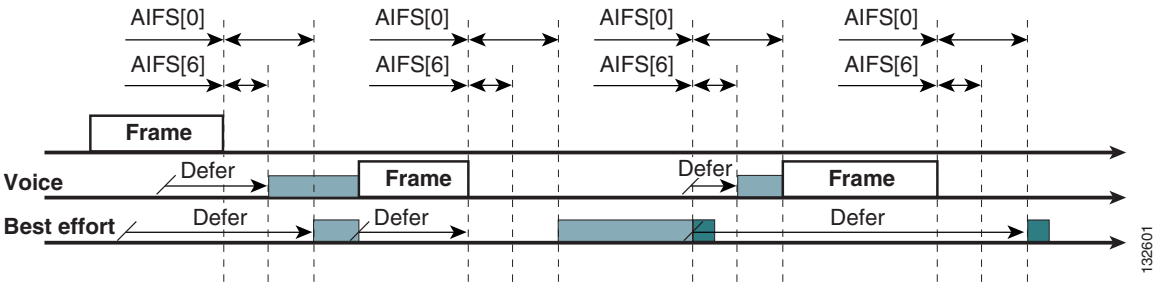

Table 5-3 WMM Client Parameters

AC	CWmin	CWmax	AIFSN	TXOP Limit (802.11b)	TXOP Limit (802.11a/g)
AC_BK	aCWmin	aCWmax	7	0	0
AC_BE	aCWmin	4*(aCQmin+1)-1	3	0	0

Table 5-3 WMM Client Parameters (continued)

AC_VI	(aCWmin+1)/2-1	aCWmin	1	6.016ms	3.008ms
AC_VO	(aCWmin+1)/4-1	(aCWmin+1)/2-1	1	3.264ms	1.504ms

Table 5-4 WMM AP Parameters

Access Category	CWmin	CWmax	AIFSN	TXOP Limit (802.11b)	TXOP Limit (802.11a/g)
AC_BK	aCWmin	aCWmax	7	0	0
AC_BE	aCWmin	4*(aCQmin+1)-1	3	0	0
AC_VI	(aCWmin+1)/2-1	aCWmin	2	6.016ms	3.008ms
AC_VO	(aCWmin+1)/4-1	(aCWmin+1)/2-1	2	3.264ms	1.504ms

The overall impact of the different AIFS, CWmin, and CWmax values is difficult to illustrate in timing diagrams because their impact is more statistical in nature. It is easier to compare the AIFS and the size of the random backoff windows, as shown in Figure 5-10.

Figure 5-10 AIFS and CWmin for Different Access Categories

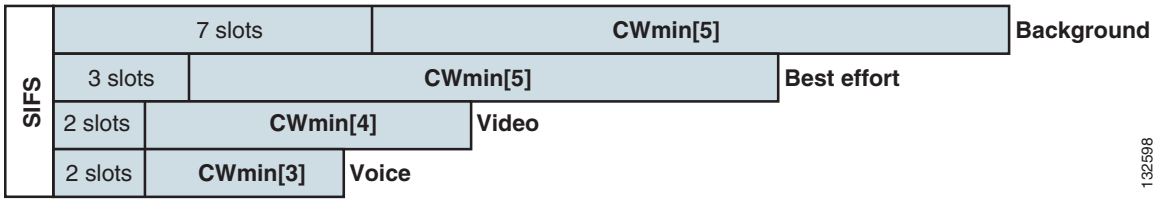

When comparing voice and background frames as examples, these traffic categories have CWmin values of 3 (7) and 5 (31), and AIFS of 2 and 7, respectively. This an average delay of 5 slot times before sending a voice frame, and an average of 22 slot times for background frame. Therefore, voice frames are statistically much more likely to be sent before background frames.

U-APSD

Unscheduled automatic power-save delivery (U-APSD) is a feature that has two key benefits:

- The primary benefit is the saving of WLAN client power, by allowing the transmission of frames from the WLAN client to trigger the forwarding of data frames for a client that has been buffered at the AP for power saving purposes.

  The client remains listening to the AP until it receives a frame from the AP with an end of service period (EOSP) bit set. This tells the client that it can now go back into its power save mode. This triggering mechanism is considered a more efficient use of client power than the regular listening for beacons method, at a period controlled by the delivery traffic indication message (DTIM) interval, as the latency and jitter requirements of voice are such that a WVoIP client would either not save power during a call, resulting in reduced talk times, or use a short DTIM interval, resulting

in reduced standby times. The use of U-APSD allows the use of long DTIM intervals to maximize standby time without sacrificing call quality. The U-APSD feature can be applied across access categories; U-APSD can be applied to the voice ACs in the AP, but the other ACs can still use the standard power save feature.

- The secondary benefit of this feature is increased call capacity. The coupling of transmission buffered data frames from the AP with the triggering data frame from the WLAN client allows the frames from the AP to be sent without the accompanying interframe spacing and random backoff, thereby reducing the contention experience by call.

Figure 5-11 shows an example frame exchange for the standard 802.11 power save delivery process. The client in power save mode first detects that there is data waiting for it at the AP via the presence of the TIM in the AP beacon. The client must power-save poll (PS-Poll) the AP to retrieve that data. If the data sent to the client requires more than one frame to be sent, the AP indicates this in the sent data frame. This process requires the client to continue sending power save polls to the AP until all the buffered data is retrieved by the client.

This presents two major problems. The first is that it is quite inefficient, requiring the PS-polls, as well as the normal data exchange, to go through the standard access delays associated with DCF. The second issue, being more critical to voice traffic, is that retrieving the buffered data is dependent on the DTIM, which is a multiple of the beacon interval. Standard beacon intervals are 100mS and the DTIM interval can be integer multiples of this. This introduces a level of jitter that is generally unacceptable for voice calls, and voice handsets switch from power save mode to full transmit and receive operation when a voice call is in progress. This gives acceptable voice quality but reduces battery life.

Figure 5-11 Standard Client Power Save

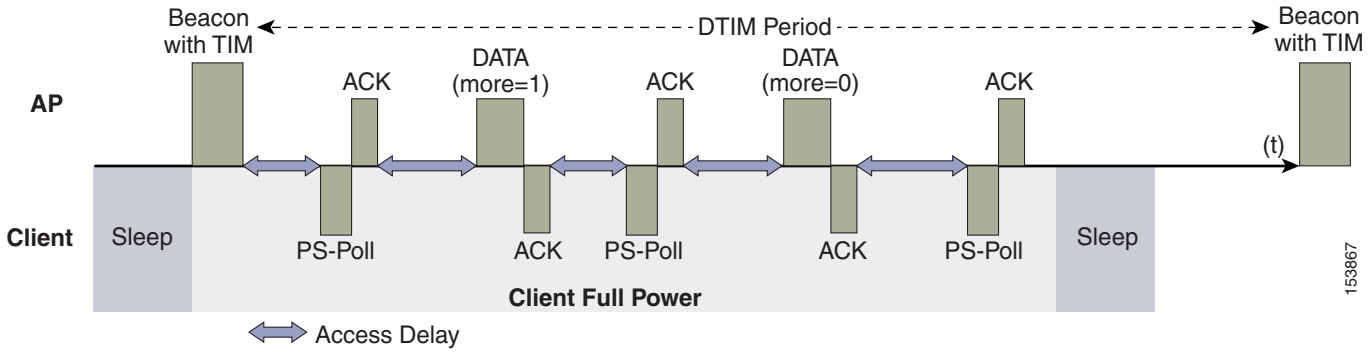

Figure 5-12 shows an example of traffic flows with U-APSD. In this case, the trigger for retrieving traffic is the client sending traffic to the AP. The AP, when acknowledging the frame, tells the client that data is queued for it, and that is should stay on. The AP then sends data to the client typically as a TXOP burst where only the first frame has the EDCF access delay. All subsequent frames are then sent directly after the acknowledgment frame.

This approach overcomes both of the disadvantages of the previous scheme in that it is much more efficient. The timing of the polling is controlled via the client traffic, which in the case of voice is symmetric, so if the client is sending a frame every 20mSec, it would be expecting to receive a frame every 20mSec as well. This would introduce a maximum jitter of 20mSec, rather than an n * 100mSec jitter.

Figure 5-12 U-APSD

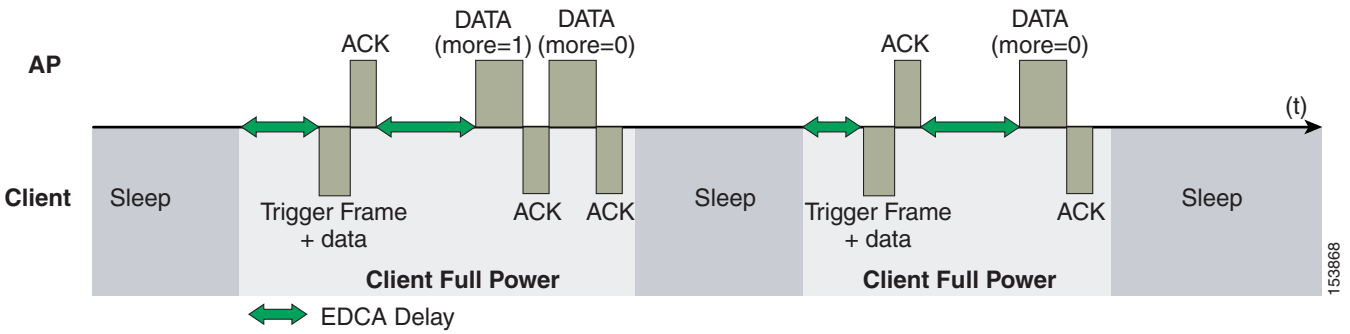

TSpec Admission Control

Traffic Specification (TSpec) allows an 802.11e client to signal to the AP its traffic requirements. In the 802.11e MAC definition, there are two mechanisms to provide prioritized access. These are the contention-based EDCA option and the controlled access option provided by the transmit opportunity (TXOP).

When describing TSpec features where a client can specify its traffic characteristics, it is easy to assume that this would automatically result in the use of the controlled access mechanism, and have the client granted a specific TXOP to match the TSpec request.

This does not have to be the case; a TSpec request can be used to control the use of the various access categories (ACs) in EDCA. Before a client can send traffic of a certain priority type, it must have requested to do so via the TSpec mechanism. For example, a WLAN client device wanting to use the voice AC must first make a request for use of that AC. Whether AC use is controlled by TSpec requests or is openly configurable can be controlled by TSpec requests, but best-effort and background ACs can be open for use without a TSpec request.

The use of EDCA ACs, rather than the HCCA, to meet TSpec requests is possible in many cases because the traffic parameters are sufficiently simple to allow them to be met by allocating capacity, rather than creating a specific TXOP to meet the application requirements.

The TSpec admission control to an AC acts in a manner similar to Cisco CallManager in that it knows how much capacity is available at a branch (AP), and does not allow additional calls when that capacity is consumed. This is done without the CallManager having to specifically allocate resources in the path of the VoIP call.

Note The Cisco 7920 WVoIP handset does not support TSpec admission control.

Add Traffic Stream

The Add Traffic Stream (ADDTS) function is how a WLAN client performs an admission request to an AP. Signalling its TSpec to the AP.

An admission request is in one of two forms:

- ADDTS action frame—This happens when a phone call is originated or terminated by a client associated to the AP. The ADDTS contains TSpec and might contain a traffic stream rate set (TSRS) IE (CCXv4 clients).

- Association and re-association message—The association message might contain one or more TSpecs and one TSRS IE if the STA wants to establish traffic stream as part of the association. The re-association message might contain one or more TSpecs and one TSRS IE if a STA roams to another AP.

The ADDTS contains the TSpec element that describes the traffic request (see Table 5-5). Apart from key data describing the traffic requirements, such as data rates and frame sizes, the TSpec element also tells the AP the minimum physical rate that the client device will use. This allows the calculation of how much time that station can potentially consume in sending and receiving in this TSpec, and therefore allowing the AP to calculate whether it has the resources to meet the TSpec.

TSpec admission control is used by the WLAN client (target clients are VoIP handsets) when a call is initiated and during a roam request. During a roam, the TSpec request is appended to the re-association request.

Table 5-5 WMM TSpec Element Field

Field	Value
Element ID	221
Length	6+55=61
OUI	00:50:f2(hex)
OUI	Type2
OUI	Subtype2
Version	1
TS Info	• Traffic stream ID, which combined with the addressing of the frame containing the TSpec element, uniquely identifies the Traffic for which a request is being made. • 802.1D priority information, and is the same value used in QoS data frames associated with this traffic stream. • Traffic is for upstream, downstream, or bi-directional traffic. • Power save is traditional or U-APSD.
Nominal MSDU Size	Size of MSDU, if fixed Nominal Size if variable
Maximum MSDU Size	-
Minimum Service Interval	-
Maximum Service Interval	-
Inactivity Interval	-
Suspension Interval	-
Service Start	Time -
Minimum Data	Rate -
Mean Data Rate	Average data rate, in units of bits per second; does not include overhead.
Peak Data Rate	-
Maximum Burst Size	-
Delay Bound	-

Table 5-5 WMM TSpec Element Field (continued)

Minimum PHY Rate	The minimum 802.11 PHY rate that is used.
Surplus Bandwidth Allowance	-
Medium Time	-

Sample TSpec Decode

```
Vendor Specific: WME            Tag Number: 221 (Vendor Specific)
Tag length: 61
Tag interpretation: WME TSPEC: type 2, subtype 2, version 1
Tag interpretation: WME TS Info: Priority 4 (Controlled Load) (Video), Contention-based
access set, Bi-directional
Tag interpretation: WME TSPEC: Fixed MSDU Size 5632
Tag interpretation: WME TSPEC: Maximum MSDU Size 56448
Tag interpretation: WME TSPEC: Minimum Service Interval 5
Tag interpretation: WME TSPEC: Maximum Service Interval 0
Tag interpretation: WME TSPEC: Inactivity Interval 0
Tag interpretation: WME TSPEC: Service Start Time 4294967040
Tag interpretation: WME TSPEC: Minimum Data Rate 255
Tag interpretation: WME TSPEC: Mean Data Rate 40960
Tag interpretation: WME TSPEC: Maximum Burst Size 40960
Tag interpretation: WME TSPEC: Minimum PHY Rate 40960
Tag interpretation: WME TSPEC: Peak Data Rate 0
Tag interpretation: WME TSPEC: Delay Bound 0
Tag interpretation: WME TSPEC: Medium Time 23437
```

QoS Advanced Features for WLAN Infrastructure

Cisco Centralized WLAN Architecture has multiple QoS features, in addition to WMM support. Primary among these is the QoS profiles in the WLC. Four QoS profiles can be configured: platinum, gold, silver, and bronze, as shown in Figure 5-13.

Figure 5-13 WLC QoS Profiles

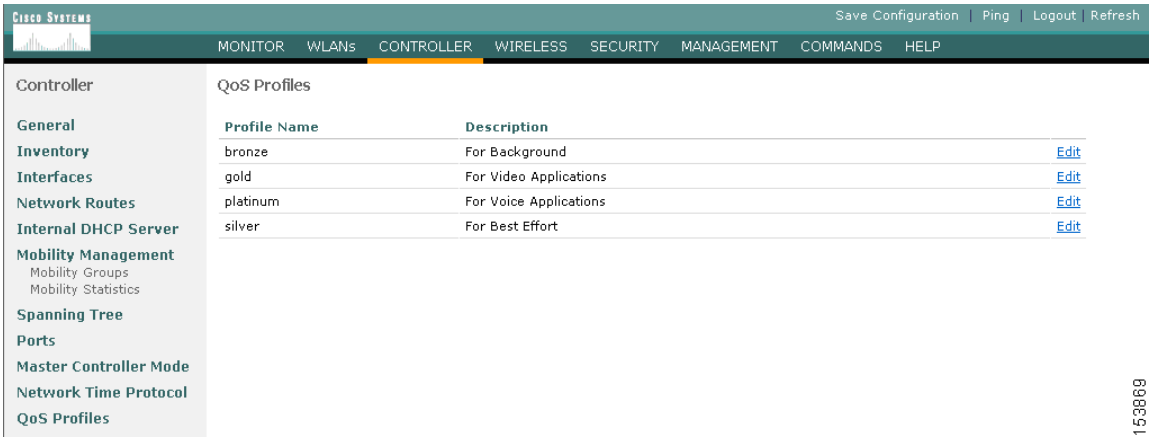

Each of these profiles (see Figure 5-14) allows the configuration of bandwidth contracts, RF usage control, and the maximum 802.1p classification allowed. It is generally recommended that these settings be left at their default values, and that the 802.11 WMM features be used to provide differentiated services.

Figure 5-14 QoS Profile Options

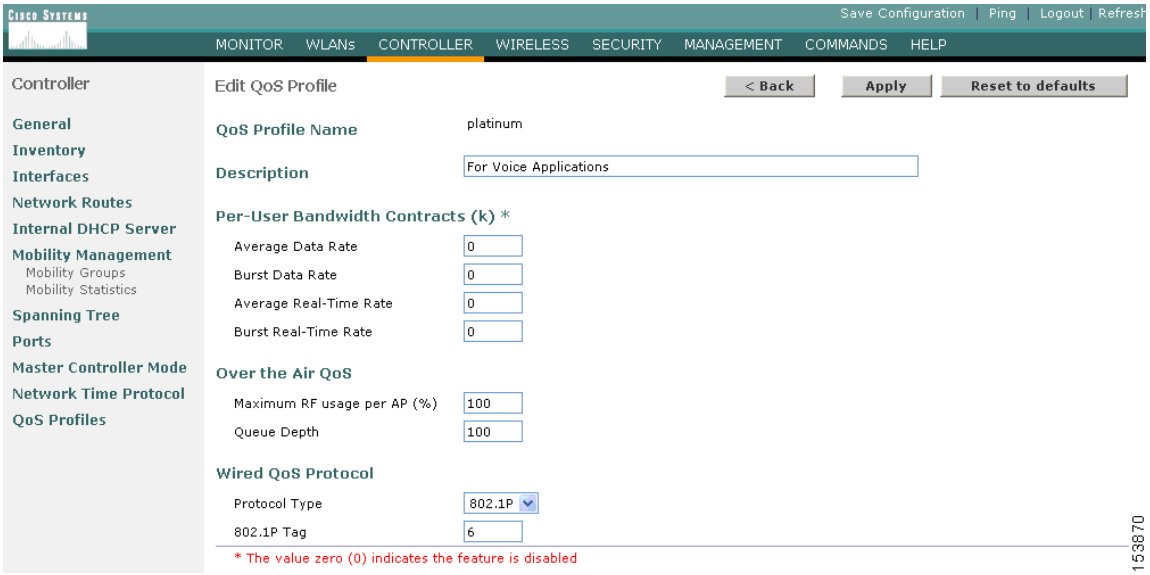

The WLAN can be configured with different default QoS profiles, as shown in Figure 5-15. Each of the profiles (platinum, gold, silver, and bronze) are annotated with their typical use. In addition, clients can be assigned a QoS profile based on their identity, through AAA.

For a typical enterprise, WLAN deployment parameters, such as per-user bandwidth contracts and over the air QoS should be left at their default values, and standard QoS tools, such as WMM and wired QoS, should be used to provide optimum QoS to clients.

Figure 5-15 WLAN QoS Profile Settings

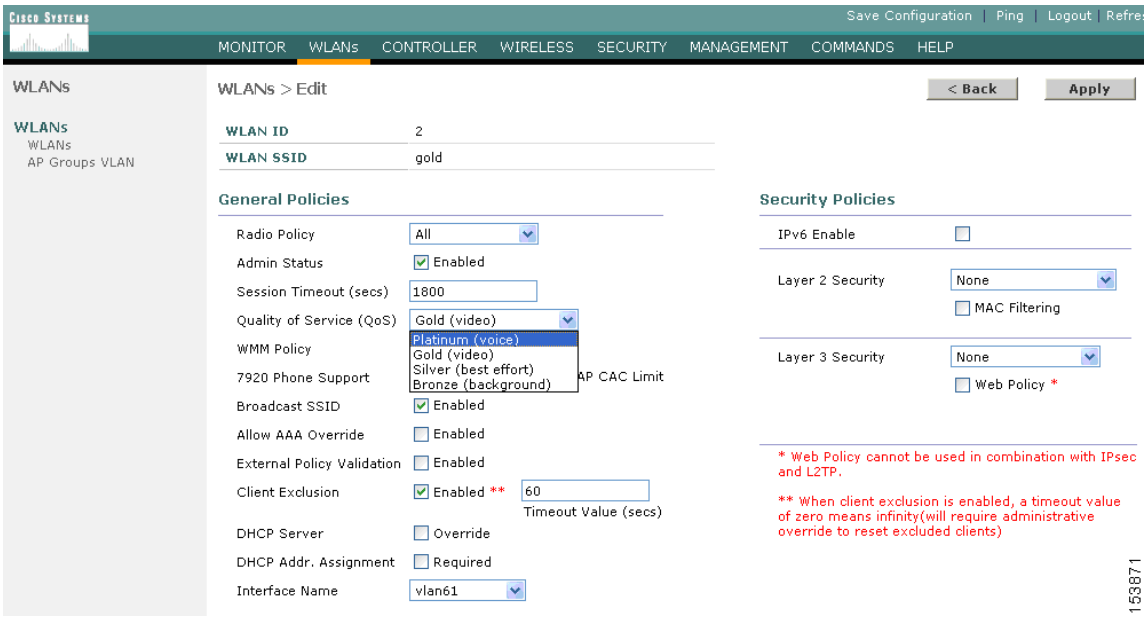

In addition to the QoS profiles, the WMM policy per WLAN can also be controlled, as shown in Figure 5-16. The three WMM options are disabled, allowed, or required. Disabled means that the WLAN does not advertise WMM capabilities, or allow WMM negotiations, Allowed means that the WLAN does allow WMM and non-WMM clients, and required means that only WMM-enabled clients can be associated with this WLAN.

Figure 5-16 WLAN WMM Policy

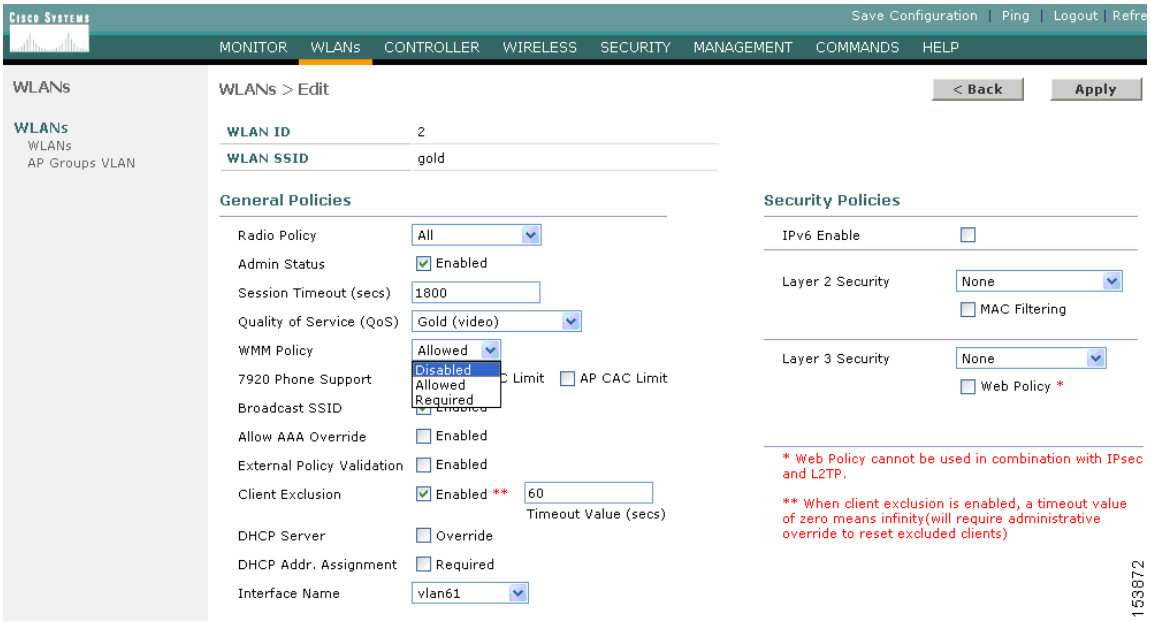

IP Phones

Figure 5-17 shows the basic QBSS information element (IE) advertised by a Cisco AP. The Load field indicates the portion of available bandwidth currently used to transport data on that AP.

Figure 5-17 QBSS Information Element

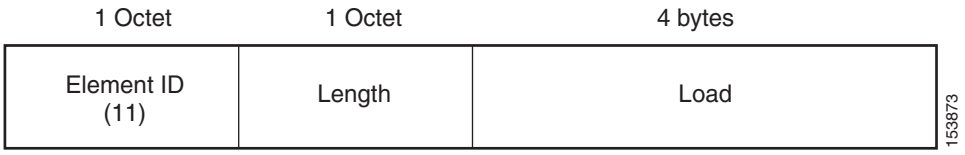

There are actually three QBSS IEs that need to be supported in certain situations:

- Old QBSS (Draft 6 (pre-standard))
- New QBSS (Draft 13 802.11e (standard))
- New distributed CAC load IE (a Cisco IE)

This QBSS depends on the WMM and 7920 settings on the WLAN.

7920 phone support, shown in Figure 5-18, is a component of the WLC WLAN configuration that enables the AP to include the appropriate QBSS element in its beacons. WLAN clients with QoS requirements use these advertised QoS parameters to determine the best AP with which to associate.

Figure 5-18 7920 Phone Support

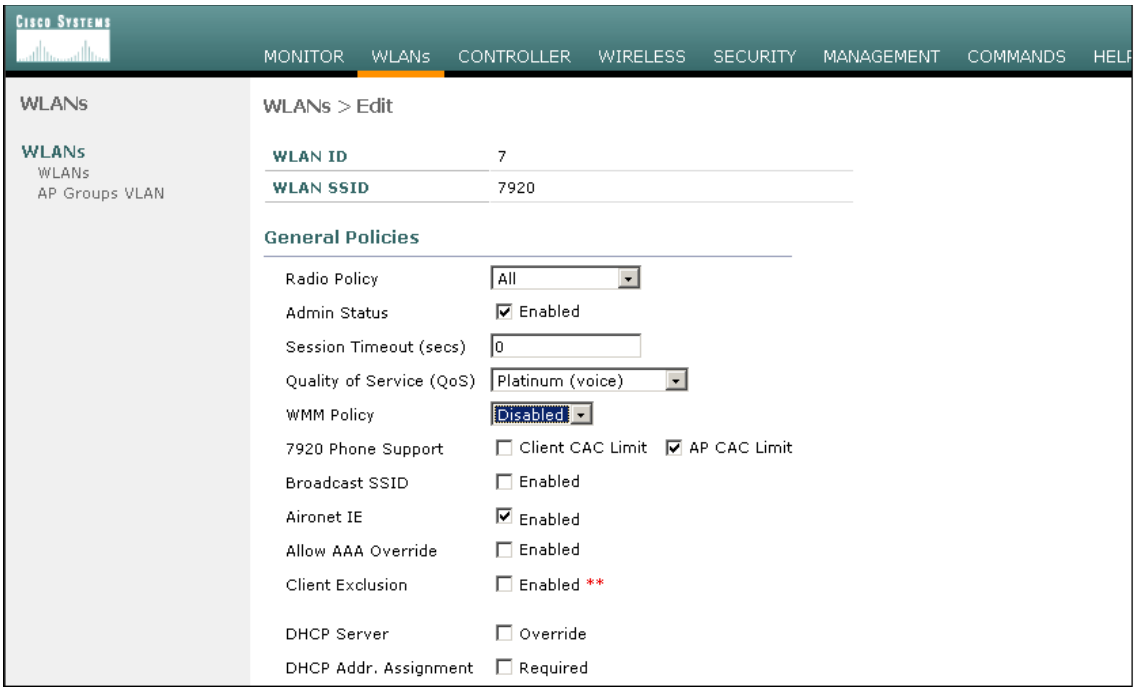

The WLC provides 7920 support through the client CAC limit, or AP CAC limit. These features provide the following:

- Client CAC limit—The 7920 uses a call admission control setting that is set on the client. This supports legacy 7920 code -pre 2.01.

- AP CAC limit—The 7920 uses call admission control settings learned from WLAN advertisement.

The various combinations of WMM, client CAC limit, and AP CAC limit result in different QBSS IEs being sent:

- If WMM only is enabled, IE number 2 (802.11e standard) QBSS Load IE is sent out in the beacons and probe responses.
- If 7920 client CAC limit is to be supported, IE number 1 (the pre-standard QBSS IE) is sent out in the beacons and probe responses on the border gateway (bg) radios.
- If 7920 AP CAC limit is to be supported, the number 3 QBSS IE is sent in the beacons and probe responses for border gateway (bg) radios.

Note The various QBSS IEs use the same ID, and therefore the three QBSSs are mutually exclusive. For example, the beacons and probe responses can contain only one QBSS IE.

Setting the Admission Control Parameters

Figure 5-19 shows an example configuration screen for setting the voice parameters on the controller. The admission control parameters consist of the maximum RF capacity that a radio can have and still accept the initiation of a VoIP call through a normal ADDTS request. The reserved roaming bandwidth is how much capacity has been set aside to be able to respond to ADDTS requests during association or re-association, which are WVoIP clients with calls in progress that are trying to roam to that AP.

When checked, the Gratuitous Probe Response checkbox causes APs to send a probe response at a regular interval (default 10mSec) to assist certain WVoIP phones in making roaming decisions. The use of the probe response is considered more efficient and less disruptive than increasing the beacon rate of APs.

The Traffic Stream Metrics Collection option determines if data is collected on voice or video calls for use by the WCS.

Note Call admission control is performed only for voice and video QoS profiles.

Figure 5-19 Configuring Voice Parameters

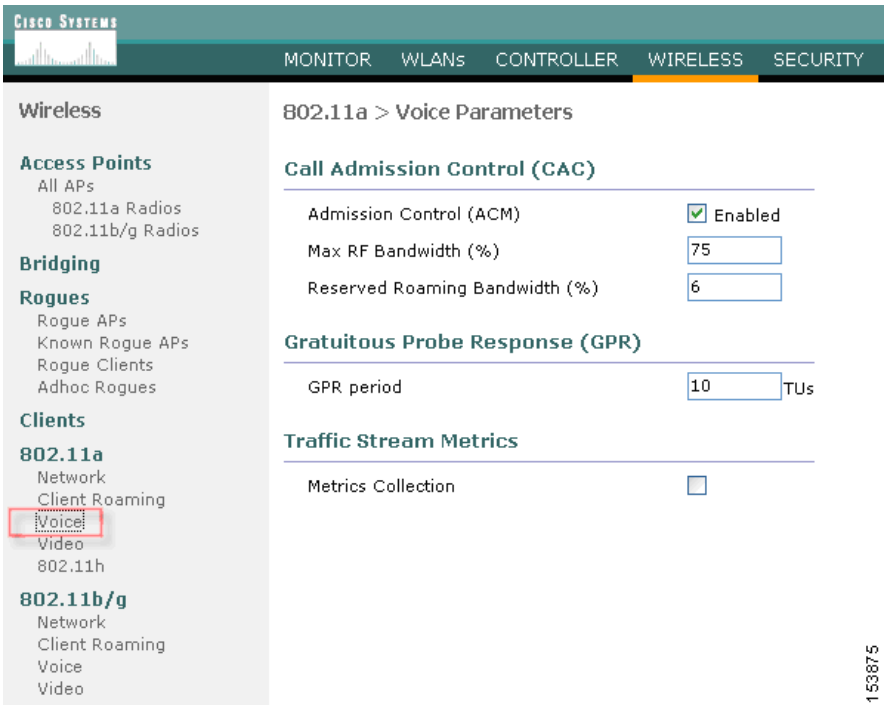

Impact of TSpec Admission Control

The purpose of TSpec admission control is not to deny clients access to the WLAN; it is to protect the high priority resources. Therefore, a client that has not used TSpec admission control does not have its traffic blocked; it simply has its traffic re-classified if it tries to send (which it should not do if the client is transmitting WMM compliant-traffic in a protected AC).

Table 5-6 and Table 5-7 describe the impact on classification if Access Control is enabled and dependent on whether a traffic stream has been established.

Table 5-6 Upstream Traffic

	Traffic Stream established	No Traffic Stream
No admission control	No change in behavior; the packets go into the network as they do today – UP is limited to max = WLAN QoS setting.	No change in behavior; the packets go into the network as they do today – UP is limited to max = WLAN QoS setting.
Admission control	No change in behavior; the packets go into the network as they do today – UP is limited to max = WLAN QoS setting.	Packets are remarked to BE (both CoS and DSCP) before they enter the network for WMM clients. For non-WMM clients, packets are sent with WLAN QoS.

Table 5-7 Downstream Traffic

	Traffic Stream established	No Traffic Stream
No admission control	No change	No change
Admission control	No change	Remark UP to BE for WMM client. For non-WMM clients, use WLAN QoS.

802.11e, 802.1p, and DSCP Mapping

WLAN data in a centralized WLAN deployment is tunneled via LWAPP. To maintain the QoS classification that has been applied to data traffic, a process transferring or applying classification is required.

For example, when WMM classified traffic is sent by a WLAN client, it has an 802.1p classification in its frame. The AP needs to translate this classification into a DSCP value for the LWAPP packet carrying the frame to ensure that the packet is treated with the appropriate priority on its way to the WLC. A similar process needs to occur on the WLC for LWAPP packets going to the AP.

A mechanism to classify traffic from non-WMM clients is also required, so that their LWAPP packets can also be given an appropriate DSCP classification by the AP and the WLC.

Figure 5-20 shows the various classification mechanisms in the LWAPP WLAN network. The multiple classification mechanisms and client capabilities require multiple strategies:

- LWAPP control frames require prioritization, and LWAPP control frames are marked with a DSCP classification of CS6.
- WMM-enabled clients have the classification of their frames mapped to a corresponding DSCP classification for LWAPP packets to the WLC. This mapping follows the standard IEEE CoS to DSCP mapping, with the exception of the changes necessary for AVVID compliance. This DSCP value is translated at the WLC to a CoS value on the wired interfaces.
- Non-WMM clients have the DSCP of their LWAPP tunnel set to match the default QoS profile for that WLAN. For example, the QoS profile for a WLAN supporting 7920 phones would be set to platinum, resulting in a DSCP classification of EF for data frames packets from that AP WLAN.

- LWAPP data packets from the WLC have a DSCP classification that is determined by the DSCP of the wired data packets sent to the WLC. The WMM classification used when sending frames from the AP to a WMM client is determined by the AP table converting DSCP to WMM classifications.

Figure 5-20 WMM, DSCP, and 802.1p Relationship

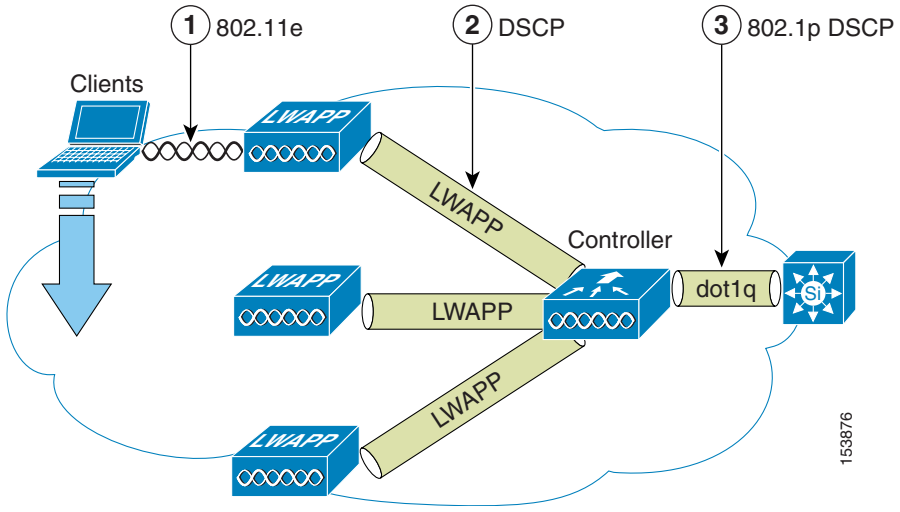

AVVID Priority Mapping

The LWAPP AP and WLC perform AVVID conversion, so that WMM values as shown in Table 5-8 are mapped to the appropriate AVVID DSCP values, rather than the IEEE values.

Table 5-8 Access Point QoS Translation Values

AVVID 802.1p UP-Based Traffic Type	AVVID IP DSCP	AVVID 802.1p UP	IEEE 802.11e UP
Network control	-	7	-
Inter-network control (LWAPP control, 802.11 management)	48	6	7
Voice	46 (EF)	5	6
Video	34 (AF41)	4	5
Voice control	26 (AF31)	3	4
Background (Gold)	18 (AF21)	2	2
Background (Gold)	20 (AF22)	2	2
Background (Gold)	22 (AF23)	2	2
Background (Silver)	10 (AF11)	1	1
Background (Silver)	12 (AF12)	1	1
Background (Silver)	14 (AF13)	1	1
Best Effort	0 (BE)	0	0, 3
Background	2	0	1

Table 5-8 Access Point QoS Translation Values

Background	4	0	1
Background	6	0	1

Deploying QoS Features Cisco on LWAPP-based APs

When deploying WLAN QoS on the APs, consider the following:

- In the absence of Layer 2 classification (802.1p) information, the WLC and the APs depend on Layer 3 classification (DSCP) information. This DSCP value is subject to modification by intermediate routers and therefore the Layer 2 classification received by the destination might not reflect the Layer 2 classification marked by the source of the LWAPP traffic.
- The APs no longer use NULL VLAN ID. As a consequence, L2 LWAPP does not effectively support QoS because the AP does not send the 802.1p/Q tags, and in L2 LWAPP there is no outer DSCP to fall back on.
- APs do not classify packets; they prioritize packets based on CoS value, or WLAN.
- APs carry out EDCF like queuing on the radio egress port only.
- APs only do FIFO queueing on the Ethernet egress port.

QoS and the H-REAP

For WLANs that have data traffic forwarded to the WLC, behavior is same as non H-REAP APs.

For locally switched WLANs with WMM traffic, the AP marks the dot1p value in the dot1q VLAN tag for upstream traffic. This occurs only on tagged VLANs; that is, not native VLANs. For downstream traffic for locally switched WLANs, the H-REAP uses the incoming dot1q tag from the Ethernet side and marks the WMM values on the wireless side. The WLAN QoS profile is applied both for upstream and downstream packets; that is, for downstream if an 802.1p value that is higher than the default WLAN value is received, the default WLAN value is used. For upstream, if the client sends an WMM value that is higher than the default WLAN value, the default WLAN value is used.

For non-WMM traffic, there is no CoS marking on the client frames from the AP.

Guidelines for Deploying Wireless QoS

The same rules for deploying QoS in a wired network apply to deploying QoS in a wireless network. The first and most important guideline in QoS deployment is to know your traffic. Know your protocols, your application's sensitivity to delay, and traffic bandwidth. QoS does not create additional bandwidth, it simply gives more control of where the bandwidth is allocated.

Throughput

An important consideration when deploying 802.11 QoS is to understand the offered traffic, not only in terms of bit rate, but also in terms of frame size, because 802.11 throughput is sensitive to the frame size of the offered traffic.

Table 5-9 shows the impact that frame size has on throughput: as packet size decreases, so does throughput. For example, if an application offering traffic at a rate of 3Mbps is deployed on an 11Mbps 802.11b network, but uses an average frame size of 300 bytes, no QoS setting on the AP allows the application to achieve its throughput requirements. This is because 802.11b cannot support the required throughput for that throughput and frame size combination. The same amount of offered traffic, having a frame size of 1500 bytes, does not have this issue.

Table 5-9 Throughput Compared to Frame Size

	300	600	900	1200	1500	Frame Size (Bytes)
11g - 54Mbps	11.4	19.2	24.6	28.4	31.4	Throughput bps
11b - 11Mbps	2.2	3.6	4.7	5.4	6	Throughput bps

Traffic Shaping, Over the Air QoS and WMM Clients

Traffic shaping and over the air QoS are useful tools in the absence of WLAN WMM features, but they do not address the prioritization of 802.11 traffic directly. For WLANs that support WMM clients or 7920 handsets, the WLAN QoS mechanisms of these clients should be relied on; no traffic shaping or over the air QoS should be applied to these WLANs.

WLAN Voice and the Cisco 7920

The Cisco 7920 is a Cisco 802.11b VoIP handset, and its use is one of the most common reasons for deploying QoS on a WLAN.

Deploying voice over WLAN infrastructure involves more than simply providing QoS on WLAN. A voice WLAN needs to consider site survey coverage requirements, user behavior, roaming requirements and admission control. This is covered in the *Cisco Wireless IP Phone 7920 Design and Deployment Guide*, at the following URL:

http://www.cisco.com/en/US/products/hw/phones/ps379/products_implementation_design_guide_book09186a00802a029a.html

Cisco Unified Wireless Multicast Design

Introduction

This chapter describes the improvements that have been made in IP multicast forwarding and provides information on how to deploy multicast in a wireless environment. A prerequisite for using the new multicast performance functionality is that a multicast-enabled network must be configured on all routers between the controllers and the APs. To accommodate networks that do not support multicast, the controller continues to support the original unicast packet forwarding mechanism.

Overview of Multicast Forwarding

Before Cisco Unified Wireless Network Software Release 3.2, when IP multicast was enabled, the controller delivered multicast packets to WLAN clients by making copies of the multicast packets and then forwarding the packets through a unicast Lightweight Access Point Protocol (LWAPP) tunnel to each AP connected to the controller. Each multicast frame received by the controller from a VLAN on the first hop router was copied and sent over the LWAPP tunnel to each of the access points connected to it, as shown in Figure 6-1. The unicast LWAPP packet containing the multicast packet used a WLAN bitmap, which told the receiving AP which WLAN SSIDs it must forward the packet over (for example, all WLAN SSIDs associated with the incoming VLAN). When the AP received the LWAPP packet, it stripped off the outer LWAPP encapsulation and transmitted a copy to each WLAN SSID (on all radios associated to the WLAN SSID) identified in the LWAPP WLAN ID bitmask.

Note Enabling multicast packet forwarding also enables broadcast packet forwarding in either the unicast mode or multicast mode of forwarding; the WLC still blocks the ARP broadcast from the WLAN, but because IP broadcast is simply a special cast of multicast, it is forwarded.

Figure 6-1 Multicast Forwarding Mechanism in Release 3.1 and Earlier Versions

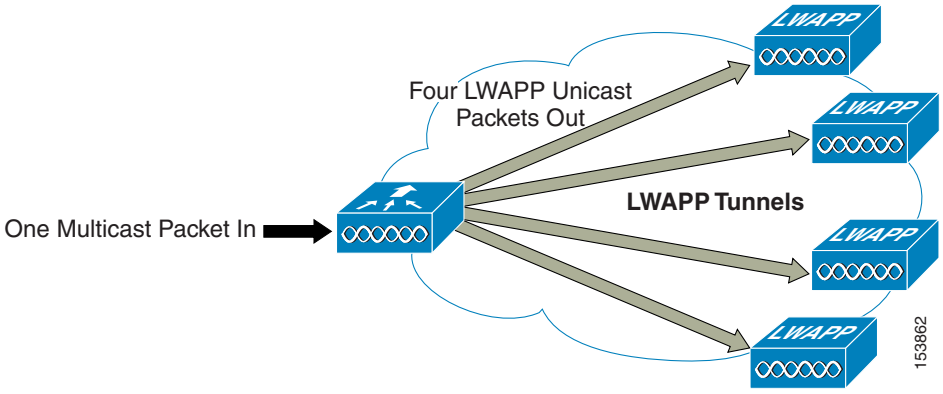

Depending on the number of APs, the controller might need to generate up to 300 copies of each multicast packet. This mechanism is inefficient, and places a large processing burden on the controller, flooding the network with a large number of duplicate unicast packets.

In Cisco Unified Wireless Network Software Release 3.2 and later releases, the multicast performance of the Cisco Unified Wireless Network has been optimized, by introducing a more efficient way of delivering multicast traffic from the controller to the access points. Instead of using unicast to deliver each multicast packet over the LWAPP tunnel to each access point, an LWAPP multicast group is used to deliver the multicast packet to each access point (see Figure 6-2). This allows the routers in the network to use standard multicast techniques to replicate and deliver multicast packets to the APs. For the LWAPP multicast group, the controller becomes the multicast source and the APs become the multicast receivers. For the multicast performance feature, the APs accept IGMP queries only from the router and multicast packets with a source IP address of the controller with which they are currently associated.

There are two important considerations to understand when enabling this feature: 1) enabling multicast packet forwarding either in unicast or multicast mode also enables broadcast packet forwarding, and 2) with multicast enabled, any kind of multicast packet received on the VLAN from the first hop router is transmitted over the wireless network, including HSRP hellos and all router EIGRP and PIM multicast packets. If you are using millisecond hellos with HSRP on the client VLAN, this could seriously degrade the WLAN throughput for clients.

If IP multicast is to be deployed and streamed across the wireless network, implement the following recommendations:

- Prevent unwanted multicast traffic from being sent on the air interface.
- Control which multicast groups are allowed by implementing multicast boundaries on the egress Layer 3 interface connecting to the VLAN or interface to the AP or bridge.
- To gain the highest AP/bridge performance for multicast traffic and data traffic, configure the WLAN APs to run at the highest possible fixed data rate. This removes the requirement for multicast to clock out at a slower rate, which can impact the range of the AP/bridge and must be taken into account in the site survey.
- If multicast reliability is a problem (seen as dropped packets), ignore the preceding recommendation and use a slower data rate (base rate) for multicast. This gives the multicast a better signal-to-noise ratio through the coding gain of the lower data rate, and can reduce the number of dropped packets.
- Test the multicast application for suitability in the WLAN environment. Determine the application and user performance effects when packet loss is higher than that seen on wired networks.

Two multicast scenarios are now described with the enhanced multicast forwarding algorithm: 1) the source of the multicast is on the wired network and streams multicast to wireless users (this is the typical scenario), and 2) a wireless user is a source of the multicast stream to both wired and wireless users.

Figure 6-2 Enhanced Multicast Forwarding Mechanism in Version 3.2

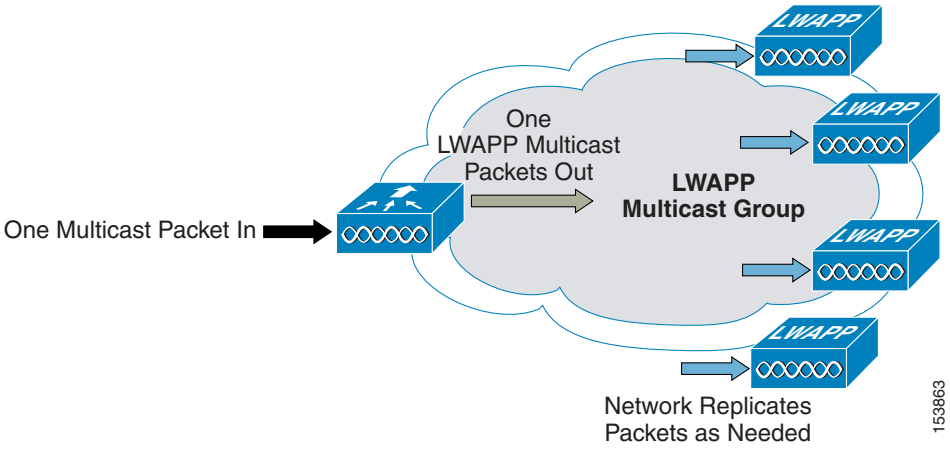

After the administrator enables multicast (multicast mode is disabled by default) and configures an LWAPP multicast group, the new multicast algorithm works in one of the following ways:

- When the source of the multicast group is on the wired LAN:
  - The LWAPP access points download the controller LWAPP multicast group address during the normal join process (at boot time) to the controller. After an access point joins a controller and downloads its configuration, the AP issues an Internet Group Management Protocol (IGMP) join request to join the controller LWAPP multicast group. This triggers the normal setup for the multicast state in the multicast-enabled routers, between the controller and APs. The source IP address for the multicast group is the controller management interface IP address, not the AP-manager IP address used for Layer 3 mode.
  - When the controller receives a multicast packet from any of the client VLANs on the first hop router, it transmits the packet to the LWAPP multicast group via the management interface at the lowest QoS level. The QoS bits for the LWAPP multicast packet are hard-coded at the lowest level and cannot be changed by the user.
  - The multicast-enabled network delivers the LWAPP multicast packet to each of the access points that have joined the LWAPP multicast group, using the normal multicast mechanisms in the routers to replicate the packet along the way, as needed, so that the multicast packet reaches all APs (see Figure 6-2). This relieves the controller from replicating the multicast packets.
  - Access points can receive other multicast packets but process only the multicast packets that come from the controller to which they are currently joined; any other copies are discarded. If more than one WLAN SSID is associated to the VLAN from where the original multicast packet was sent, the AP transmits the multicast packet over each WLAN SSID (following to the WLAN bitmap in the LWAPP header). In addition, if that WLAN SSID is on both radios (802.11g and 802.11a), both radios transmit the multicast packet on the WLAN SSID if there are clients associated with it, even if those clients did not request the multicast traffic.
- When the source of the multicast group is a wireless client:
  - The multicast packet is unicast (LWAPP-encapsulated) from the AP to the controller, like standard wireless client traffic.

– The controller makes two copies of the multicast packet. One copy is sent out the VLAN associated with the WLAN SSID on which it arrived, enabling receivers on the wired LAN to receive the multicast stream and the router to learn about the new multicast group. The second copy of the packet is LWAPP-encapsulated and is sent to the LWAPP multicast group so that wireless clients can receive the multicast stream.

Enabling the Multicast Feature

There are two tasks involved in enabling the enhanced multicast feature: 1) enabling the underlying network infrastructure for multicast operation, and 2) enabling multicast forwarding on the controllers.

Multicast-enabled Networks

A prerequisite for using this new multicast performance functionality is that a multicast-enabled network must be configured on all routers between the controllers and the APs. A multicast-enabled network allows for an efficient way to deliver a packet to many hosts across the network. IP multicast is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients. Packets are replicated as necessary at each Layer 3 point in the network. A multicast routing protocol, such as PIM, is required if there is more than one router between the controllers and the APs.

See the following URLs for more information on setting up a multicast-enabled network:

- http://www.cisco.com/go/multicast
- http://www.cisco.com/en/US/products/ps6552/prod_white_papers_list.html

Enabling Multicast Forwarding on the Controller

Enabling multicast packet forwarding also enables broadcast packet forwarding in both unicast and multicast modes.

Because of the load that replicating multicast packets places on the controller in unicast mode, IP multicast traffic through the controller is disabled by default. WLAN clients cannot receive multicast traffic when it is disabled.

If you want to enable multicast traffic to the WLAN clients, and you have a multicast-enabled network, select **multicast** under Ethernet Multicast Mode to use the method where the network replicates the packets.

If you want to enable multicast traffic to the WLAN clients, and you do not have a multicast-enabled network, select **unicast** under Ethernet Multicast Mode to use the method where the controller replicates the packets.

Commands for Enabling Ethernet Multicast Mode via the GUI

To enable the Ethernet Multicast mode using the GUI (see Figure 6-3), follow these steps:

Step 1 From the controller general web page, ensure that the LWAPP transport mode is set to Layer 3. The multicast performance feature works only in this mode.

Step 2 From the drop-down menu for the Ethernet Multicast Mode, select multicast and type in a multicast group address. In this example, we entered 239.1.1.57.

Figure 6-3 Enabling Multicast Forwarding

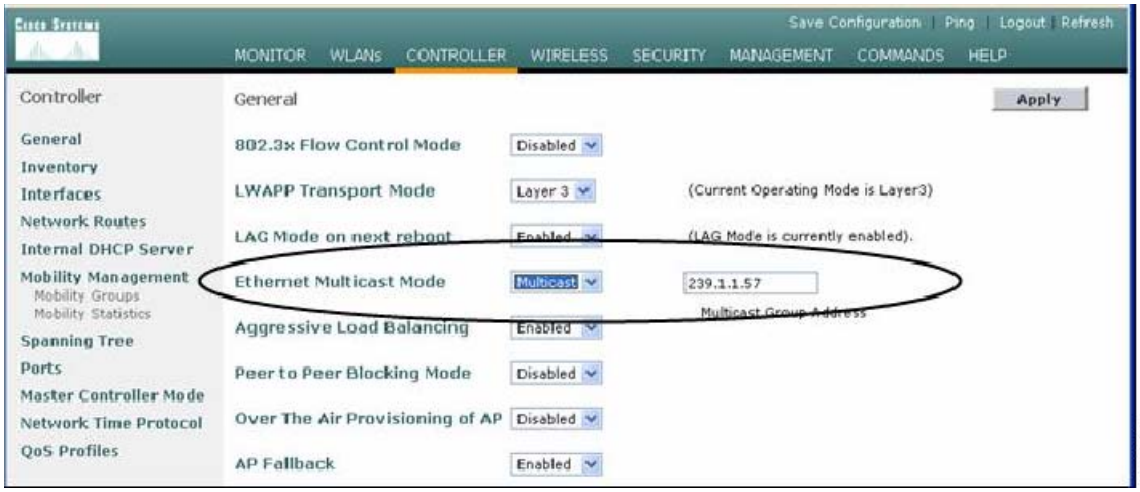

Commands for Enabling Ethernet Multicast Mode via the CLI

To enable the Ethernet Multicast mode using the CLI, follow these steps:

Step 1 From the command line, enter the **config network multicast global enable** command.

Step 2 From the command line, enter the **config network multicast mode multicast multicast-group-ip-address** command.

Use the **show network** command to verify the multicast mode on the controller and the **show lwapp mcast** command to verify the group on the AP. Other useful commands are **show ip mroute** and **show ip igmp membership** on the routers.

Note Do not confuse these commands with the multicast appliance mode on the port configuration. The 4400 controllers and WiSM use the global mode to enable multicast on their ports.

Multicast Deployment Considerations

LWAPP Multicast Reserved Ports and Addresses

The controller blocks all multicast packets sent to any multicast group that have a destination port of 12222 through 12224. Additionally, all packets with a multicast group address equal to the controller LWAPP multicast group address are blocked at the controller. This prevents fragmented LWAPP-encapsulated packets from another controller being retransmitted (see Fragmentation and LWAPP Multicast Packets, page 6-6 for more information).

Ensure that the multicast applications in your network do not use these reserved ports or LWAPP multicast group addresses.

Recommendations for Choosing an LWAPP Multicast Address

Cisco recommends that you assign multicast addresses from the administratively-scoped block 239/8. IANA has reserved the range of 239.0.0.0-239.255.255.255 as administratively-scoped addresses for use in private multicast domains. These addresses are similar in nature to the reserved Private IP unicast ranges (such as 10.0.0.0/8) that are defined in RFC 1918. Network administrators are free to use the multicast addresses in this range inside of their domain without fear of conflicting with others elsewhere in the Internet. This administrative or private address space should be used within the enterprise and blocked from leaving or entering the Autonomous System (AS).

Note You can assign any multicast address to the LWAPP multicast group, including the reserved link local multicast addresses used by OSPF, EIGRP, PIM, HSRP, and other multicast protocols, but Cisco does not recommend assigning reserved multicast addresses. Doing so would impact the application using the reserved multicast address by making it process the extra packets from the controller. This also impacts the performance of the controller by making it process the multicast packets of the application. It is much more efficient for an interface to drop packets not destined for that host. By not using the reserved address space, packets can be dropped at the interface instead of being forwarded for processing by the controller.

Cisco recommends that enterprise network administrators further subdivide this address range into smaller geographical administrative scopes within the enterprise network to limit the scope of particular multicast applications. This prevents high-rate multicast traffic from leaving a campus (where bandwidth is plentiful) and congesting the WAN links. It also allows for efficient filtering of the high bandwidth multicast to prevent it from reaching the controller and the wireless network.

Note Do not use the 239.0.0.X or the 239.128.0.X address ranges. Addresses in these ranges overlap with the link local MAC addresses and flood all switch ports even with IGMP snooping enabled. For more information on overlapping multicast MAC addresses, see the following URL: http://www.cisco.com/en/US/products/ps6552/products_white_paper09186a00800d6b5e.shtml#xtocid9

For more information on multicast address guidelines, see *Guidelines for Enterprise Multicast Address Allocation* at the following URL:

http://www.cisco.com/application/pdf/en/us/guest/products/ps6592/c1244/cdccont_0900aecd80310d68.pdf

Fragmentation and LWAPP Multicast Packets

When a controller receives a multicast packet, it LWAPP-encapsulates it using the LWAPP multicast group as a destination address and then forwards it to the APs via the management interface (source address). If the packet exceeds the MTU of the link, the controller fragments the packet into two packets and sends out both packets to the LWAPP multicast group. If another controller were to receive this LWAPP-encapsulated multicast packet via the wired network, it could re-encapsulate it, treating it as a normal multicast packet, and would then forward it to its APs.

There are two options to prevent this from happening, either of which is effective by itself: 1) you can assign all controllers to the same LWAPP multicast group address, or 2) you can apply standard multicast filtering techniques to ensure that LWAPP-encapsulated multicast packets do not reach any other controller. Table 6-1 provides the pros and cons of these two techniques.

Table 6-1 Pros and Cons of Using the Same Multicast Group or Different Groups

Option	Pros	Cons
All controllers have the same LWAPP multicast group	No need to do any additional fragmentation protection measures	Each controller's multicast traffic is flooded throughout the network (APs will drop multicast packets that do not have a source IP address that is equal to their controller management interface).
Standard multicast techniques are used to block LWAPP multicast fragments	Can use a range of addresses, thus preventing flooding throughout the network	ACL filtering must be applied on first hop router on all VLANs configured on multicast-enabled controllers.

All Controllers Have the Same LWAPP Multicast Group

To prevent the second controller from re-transmitting these LWAPP encapsulated packets, the controllers block incoming multicast packets to the LWAPP multicast group and the LWAPP reserved ports. By blocking the reserved ports, the controller blocks the first part of a fragmented packet in an encapsulated LWAPP multicast packet. However, the second packet does not contain port numbers and can be blocked only by filtering it on the multicast group address (destination address). The controller blocks any packets where the destination address is equal to the LWAPP multicast group address that is assigned to the controller.

However, assigning every controller to the same LWAPP multicast group creates other problems, although smaller. IGMP version 1 and 2 used by the APs to join the LWAPP multicast group are Any Source Multicast (ASM) and the APs receive multicast traffic from all sources of the multicast group in the network. This means that the APs receive multicast packets from all of the controllers on the network if the controllers are configured with the same multicast group address, and no multicast boundaries have been applied. One controller's multicast traffic floods out to all of the APs across the network and every AP receives (and drops, if the source address is not equal to its controller's management address) the multicast traffic that is being received from any wireless multicast client in the entire network.

Note Cisco IOS APs (such as the 1240) use IGMPv2 while VxWorks APs (such as the 1030) use IGMPv1.

Controlling Multicast on the WLAN using Standard Multicast Techniques

In the past, the Time To Live (TTL) field in the IP multicast datagram was used for creating Auto-RP administrative boundaries using the **ttl-threshold** command. This has been superseded by the **ip multicast boundary interface mode** command, which filters IP multicast traffic and also AutoRP messages. Cisco recommends transitioning to, and using, the new command.

Normal boundary techniques should be used in your multicast-enabled network. These include using the **ip multicast boundary interface mode** command, which filters IP multicast traffic and also Auto-RP messages.

Other useful commands include the **ip multicast rate-limit interface** command. This command enforces low rates on the wireless VLANs. Without it, even if the network engineer filters the high rate multicast addresses, a low rate multicast address cannot exceed its rate.

Caution A wired client anywhere in the network can request the LWAPP multicast stream and receive it from all sources (if multicast boundaries are not applied). Multicast streams are not encrypted when they are encapsulated in the LWAPP multicast packet.

A typical example on a wireless client VLAN is given below. For more information on other multicast commands for a multicast-enabled network, see the following URL:

http://www.cisco.com/go/multicast

Filtering for multicast-enabled traffic also allows you to prevent propagation of certain worms such as the Sasser worm, which relied on the TCP and ICMP transports with multicast addresses. Blocking these types of addresses with multicast group addresses does not affect most applications because they typically use UDP or RCP for streaming.

In the following example, packets to the multicast group range 239.0.0.0 to 239.127.255.255 from any source have their packets rate limited to 128 kbps. The following example also sets up a boundary for all multicast addresses that are not in the lower administratively scoped addresses: In addition, hosts serviced by Vlan40 can join only the lower administrative groups 239.0.0.0 through 239.127.255.255.

Note The Catalyst 6500 does not support the **ip multicast rate-limit interface** command. The following example uses QoS to rate limit the multicast on a Catalyst 6500. Other Layer 3 switches, such as the Catalyst 3750, support the **ip multicast rate-limit interface** command. Use the **show policy-map interface** command to verify your QoS configuration.

```
mls qos!
class-map match-all multicast_traffic
  description Permit Low Rate Multicast Range of 239.0.0.0 to 239.127.0.0
  match access-group 101
!
policy-map multicast
 description Rate Limit Multicast traffic to 2.56mps with burst of 12800 bytes
  class multicast_traffic
   police cir 2560000 bc 12800 be 12800 conform-action transmit exceed-action drop
!
interface Vlan40
 description To Wireless Clients
 ip address 10.20.40.3 255.255.255.0
 ip pim sparse-mode
 ip multicast boundary 1
 ip igmp access-group 30
 standby 40 ip 10.20.40.1
 standby 40 preempt
 service-policy output multicast
!
access-list 1 remark Permit Low Rate Multicast Range of 239.0.0.0 to 239.127.0.0 for
multicast boundary
access-list 1 permit 239.0.0.0 0.127.255.255
!
access-list 30 remark Only Allow IGMP joins to this Multicast Group Range
access-list 30 permit 239.0.0.0 0.127.255.255
!
access-list 101 remark Permit Low Rate Multicast Range of 239.0.0.0 to 239.127.0.0 for
class-map
access-list 101 permit ip any 239.0.0.0 0.127.255.255
```

How Controller Placement Impacts Multicast Traffic and Roaming

This section describes two deployments, distributed and co-located, and how they impact roaming with multicast clients. In co-located controllers that are attached to the same VLANs, such as in a data center, the multicast streams are uninterrupted when a multicast client roams from one controller to another controller.

However, the co-located deployment creates a flat multicast network. The reason co-located controllers do not affect multicast roaming is that when the multicast stream is requested from a single multicast client on a WLAN SSID, it streams out all APs on that WLAN SSID, on all radios (802.11g and 802.11a), on all controllers, even if that access point has no multicast clients associated with it that have requested the multicast traffic. If you have more than one WLAN SSID associated to the VLAN, the AP transmits the multicast packet for each WLAN SSID. Both the unicast mode LWAPP packet and the multicast mode LWAPP packet contain a WLAN bitmap that tells the receiving AP which WLAN SSIDs it must forward the packet over. When the AP receives a packet destined to the LWAPP multicast group, it strips off the outer header and handles the original multicast packet by sending a copy to each WLAN that is identified in the LWAPP WLAN id bitmask.

The distributed deployment reduces the amount of multicast traffic on the APs because, although the WLAN SSIDs are the same, the controllers are attached to different VLANs. WLAN multicast traffic depends on a client request on the VLAN of that WLC. This means that when the multicast client roams to a new controller, the client stops receiving the multicast stream unless it was already requested by a client on that WLC VLAN, or the client makes a new IGMP request. Table 6-2 lists the advantages and disadvantages of distributed and co-located deployments.

Table 6-2 Pros and Cons of Co-located Controllers and Distributed Controllers

	Pros	Cons
All co-located controllers connected to the same VLANs (subnets)	Multicast traffic started on any client VLAN is transmitted to all APs so clients roaming to any AP receive multicast stream.	If only one client requests multicast traffic, all APs attached to all controllers receive the stream and transmit it if they have any clients associated even if those clients did not request the multicast stream.
Distributed controllers on different VLANs and subnet	Multicast streams are isolated to APs attached to the controller.	Clients must request multicast stream before they can receive it after roaming to a new controller, unless it has already been requested by a client on that WLC.

Whether a multicast client has uninterrupted streaming while roaming is best summarized by the following statement: if a multicast client roams to an AP attached to a different controller (a client moves from an access point on their anchor controller to an access point on a foreign controller), the client can receive multicast packets in only one of the following two cases:

- If the foreign controller has direct connectivity to the same VLAN that the client was originally on. In this case, the foreign controller becomes the anchor controller for the client, and the client continues to receive the multicast stream uninterruptedly because the multicast is already streaming on that VLAN.
- If the foreign controller does not have direct connectivity to the original VLAN, another client on the foreign controller that is on the same VLAN as the roaming client is already receiving the same multicast stream (for example, if a second client has already requested it).

Additional Considerations

Two areas of for additional consideration in multicast deployment are when implementing AP groups, and H-REAPs and REAPs.

AP groups allow APs on the same controller to map the same WLAN (SSID) to different VLANs, depending on the AP that a client is using. If a WLAN is roaming between APs in different groups, the multicast session would break in the same manner as if a client roams between APs connected to different WLCs on separate VLANs.

REAP and H-REAP APs allow the local termination of WLANs at the network edge rather than at the WLC, and the multicast behavior is controlled at that edge. If an H-REAP WLAN is terminated on a WLC and multicast is enabled on that WLC, multicast is delivered to that H-REAP WLAN if the LWAPP multicast group is allowed to extend to the H-REAP network location.

Note Even if the LWAPP multicast packets are not able to transit the network to the H-REAP, WLAN clients on that H-REAP are able to send IGMP joins to the network connected to the WLC.

Cisco Unified Wireless Hybrid REAP

The Cisco Unified Wireless solution uses the Lightweight Access Point Protocol (LWAPP) between lightweight access points (APs) and a WLAN controller. In a typical centralized WLAN deployment, wireless user traffic and AP control and management traffic is tunneled between the AP and controller using LWAPP. The LWAPP tunnel can be established across Layer 2 or Layer 3 topologies. In a centralized architecture, the WLAN controller is responsible for the propagation of policies, QoS, and radio resource management information to each lightweight AP. The WLAN controller is also the sole point for ingress and egress of all wireless user traffic, and it ultimately enables mobility across a wireless enterprise through the use of LWAPP and Ethernet over IP (when roaming between controllers). However, when attempting to implement a centralized controller with lightweight APs that are deployed at remote branch locations, this architecture might not be a viable solution.

Remote Edge AP

Remote edge APs (REAPs) are special purpose LWAPP-based APs that are designed to be deployed in remote (branch) locations where:

- Fewer than three APs are needed to provide adequate wireless coverage for a given location. This is often more cost-effective than deploying and managing controllers at every location, especially if there are large numbers of small remote sites requiring wireless coverage.
- Wireless users at a branch or remote location require access to local network resources in addition to communicating back to a central site, or local wireless connectivity needs to be maintained during WAN link outages.
- Limited WAN bandwidth exists between the central site and a remote location where local connectivity is required. In this scenario, it would be impractical to tunnel all of the wireless user traffic to a centralized controller only to be routed back (in standard IP packets) across a bandwidth constrained WAN link to the remote site.

REAP APs are designed to address these remote branch needs by decoupling the LWAPP control plane from the wireless data plane. This allows WLANs to be terminated locally on a Layer 2 switch while LWAPP control and management data is tunneled back to a centralized WLAN controller. In this way, the benefits of a centralized architecture are preserved. Figure 7-1 provides a high level REAP topology diagram.

Figure 7-1 High Level REAP Topology

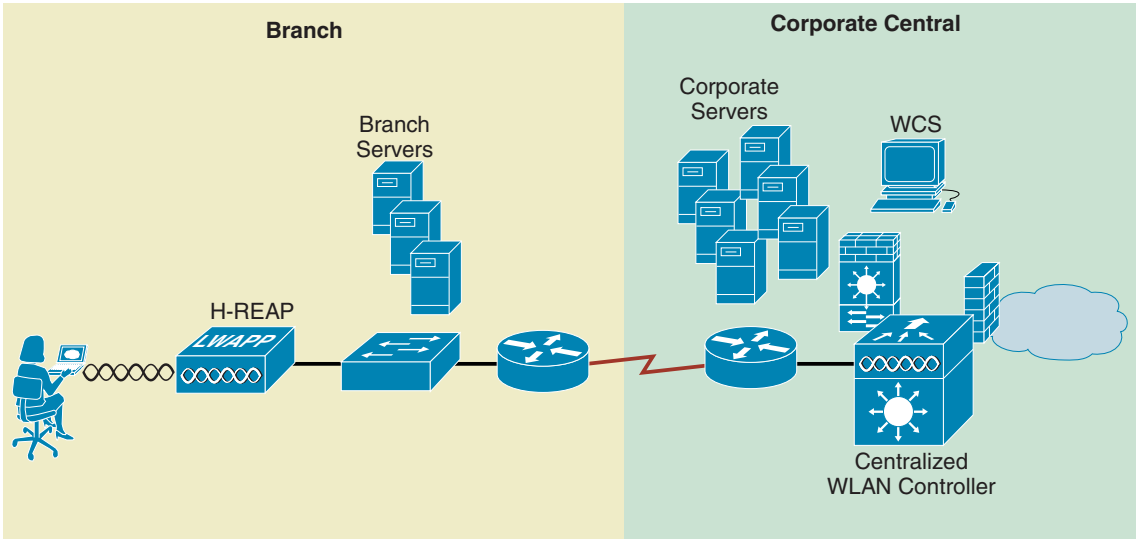

The Cisco first generation REAP, the 1030, is capable of supporting up to 16 WLANs. Although all WLANs can be locally switched, the 1030 (when in REAP mode) has some limitations compared to a standard lightweight AP that is deployed in a conventional LWAPP/ controller topology. Specifically:

- It does not support 802.1Q trunking. All WLANs terminate on a single local VLAN/subnet.
- In the event of a WAN link outage all WLANs except WLAN 1 become disabled and are no longer beaconed (if so enabled).

Cisco addressed these limitations with the introduction of a new version of REAP called Hybrid Remote Edge AP (H-REAP), which offers the ability to map WLANs to VLANs via 802.1Q trunking. Additionally, an H-REAP AP can support local switched and centrally switched WLANs concurrently. The remainder of this chapter focuses on application, features, limitations, and configuration of the H-REAP AP and, when applicable, highlights the differences between H-REAP and the older 1030 REAP platform.

Hybrid REAP

Supported Platforms

Controllers

H-REAP APs are supported by the following WLAN controller platforms with version 4.0 and later software images:

- Cisco 2000 Series
- Cisco 4400 Series
- Cisco 6500 Series (WISM)
- Cisco WLAN controller module for Integrated Service routers (ISR) Cisco C3750G-24WS

Access Points

The following IOS LWAPP APs support H-REAP functionality:

- Cisco 1130 Series
- Cisco 1240 Series

SeeAPs, page 2-10 for additional information on Cisco 1130 and 1240 series APs.

See the following URL for guidelines to convert IOS based 1130/1240 series APs to LWAPP mode of operation.:
http://www.cisco.com/en/US/products/hw/wireless/ps430/prod_technical_reference09186a00804fc3dc.html

H-REAP functionality is not supported on Cisco 1000 Series LWAPP APs. However, basic REAP functionality is still supported.

H-REAP Terminology

This section provides a summary of H-REAP terminology and definitions.

Switching Modes

Unlike the 1030 Series REAP AP, which can map wireless user traffic only to a single VLAN, H-REAP APs are capable of supporting the following switching modes concurrently, on a per-WLAN basis:

- Local Switched—Local switched WLANs map wireless user traffic to discrete VLANs via 802.1Q trunking, either to a router or switch. One or more WLANs can be mapped to the same local 802.1Q VLAN.

  A branch user who is associated to a local switched WLAN will have their traffic switched and forwarded by the on-site branch switch or router. Traffic destined off-site (to the central site) is forwarded as standard IP packets by the branch router.

  All wireless control traffic is tunneled back to the centralized controller separately via LWAPP.

- Central Switched—Central switched WLANs tunnel both the wireless user traffic and all control traffic via LWAPP to the centralized controller where the user traffic is mapped to an interface or VLAN on the controller. This is normal LWAPP mode of operation.

  The traffic of a branch user who is associated to a central switched WLAN will be tunneled directly to the centralized controller. If that user needs to communicate with computing resources within the branch (where that client is associated), their data must be forwarded as standard IP packets back across the WAN link to the branch location. Depending on the WAN link bandwidth, this might not be desirable behavior.

Operation Modes

Regardless of which switching mode is defined for a given WLAN, there is corresponding LWAPP control traffic that is sent to the controller. There are two modes of operation for an H-REAP AP:

- Connected mode—The controller is reachable. In this mode the H-REAP AP has LWAPP connectivity with its controller.
- Standalone mode—The controller is unreachable. The H-REAP has lost or failed to establish LWAPP connectivity with its controller. This would be the case when there is a WAN link outage between a branch and its central site.

Authentication Modes

The following are per-WLAN authentication modes:

- Authentication central—These are WLANs that require 802.1x, VPN, or web-based authentication services.
- Authentication local—Includes WLANs that use Open, Static, WEP or WPA PSK methods for authentication. H-REAP handles these locally, if the WAN link is down; otherwise, these are handled by the WLC.
- Authentication down—802.1x, VPN, or web authentication is unreachable because of the H-REAP AP being in standalone mode.

H-REAP States

An H-REAP WLAN, depending on its configuration and network connectivity, can be classified as being in one of the following states:

- Authentication-central / switch-central—WLAN uses centralized authentication services and user traffic is tunneled via LWAPP to the controller. Supported only when H-REAP is in Connected Mode (see Figure 7-2). 802.1X is used in this example, but other mechanisms are equally applicable.
- Authentication-central / switch-local—WLAN uses centralized authentication and user traffic is switched locally. Supported only when H-REAP is in Connected mode (see Figure 7-3). 802.1X is used in this example, but other mechanisms are equally applicable.
- Authentication-local / switch-local—WLAN uses local authentication: open, static wep or wpa psk and user traffic is switched locally at the branch. The H-REAP AP can be in connected mode or local mode (see Figure 7-4). If the AP is in connected mode, the authentication is processed by the WLC.
- Authentication-down / switch-local—An existing WLAN that requires central authentication will reject new users. Existing authenticated users continue to be switched locally. WLAN SSIDs continue to be beaconed and respond to probes. The H-REAP AP is in standalone mode because the WLC is not accessible (see Figure 7-5).

Figure 7-2 Authentication-Central / Switch-Central WLAN

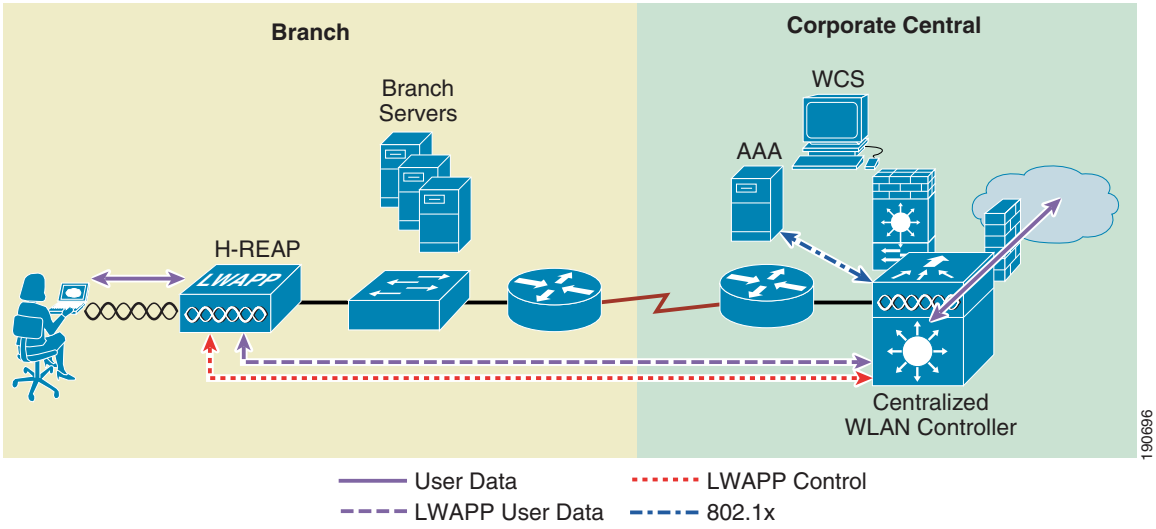

Figure 7-3 Authentication-Central / Switch-Local WLAN

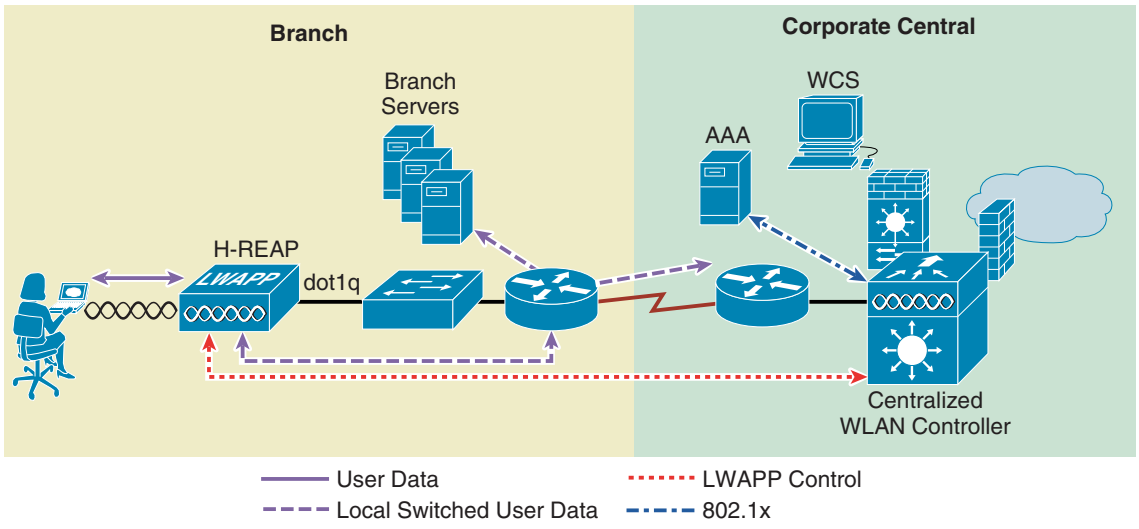

Figure 7-4 Authentication-Local / Switch-Local WLAN

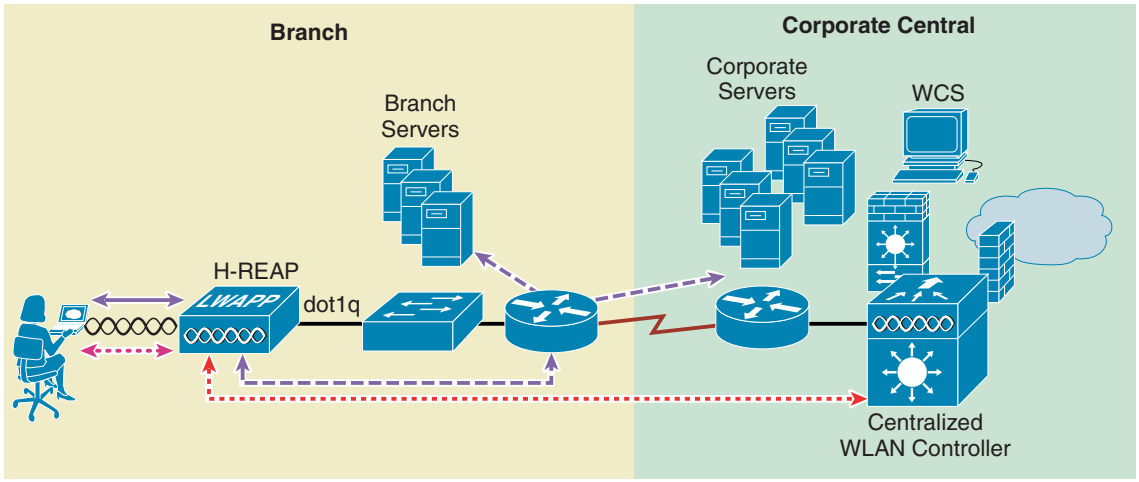

Figure 7-5 Standalone WLAN

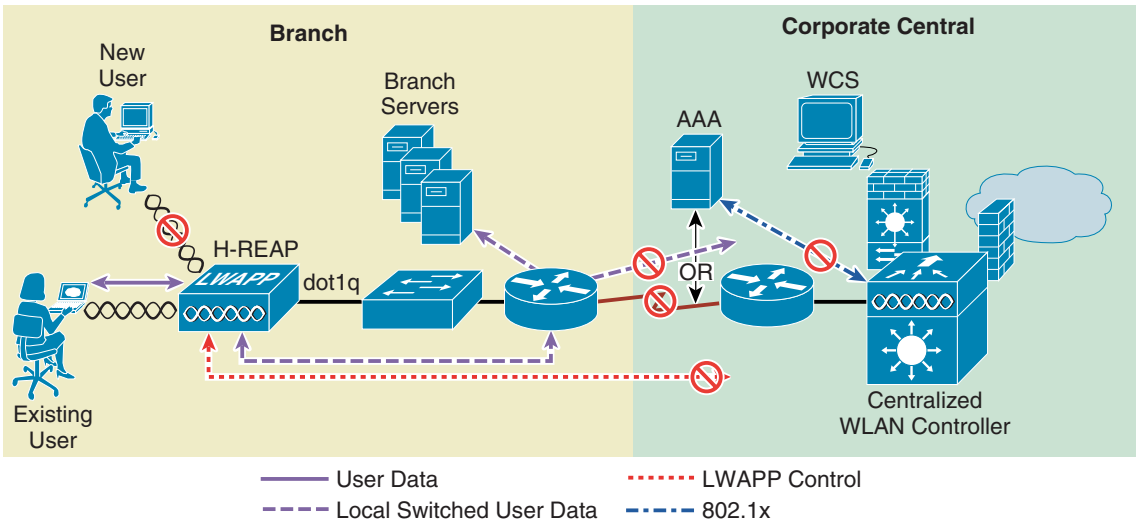

Applications

With its expanded capabilities, the H-REAP AP offers great flexibility in how it can be deployed.

Branch Wireless Connectivity

The primary goal of REAP and H-REAP is to address the wireless connectivity needs in branch locations; permitting wireless user traffic to be terminated locally rather than be tunneled across the WAN to a central controller.

Because H-REAP can map individual WLANs to specific 802.1Q VLANs, branch locations can more effectively implement segmentation, access control, and QoS policies on a per-WLAN basis.

Branch Guest Access

One of the challenging aspects of using standard REAP APs in the branch is the implementation of guest access, which is difficult to implement for the following reasons:

- All WLANs map to the same local VLAN, thereby making it difficult to differentiate and segment guest users from branch users.
- All user traffic is switched locally, guest access traffic must somehow be segmented and routed back to the central site for access control and authentication, or if local Internet access is available at the branch, both segmentation and access control must be implemented locally.

The H-REAP AP helps overcome some of these challenges with the introduction of concurrent local and central switching. In an H-REAP topology, an SSID/WLAN designated for guest access can be tunneled via LWAPP to a central controller where its corresponding interface or VLAN can be switched directly to an interface of an access control platform, such as BBSM, SSG, or Clean Access. Or the centralized controller itself can perform web authentication for the guest access WLAN. In either case, the guest user's traffic is segmented (isolated) from other branch office traffic. Figure 7-6 provides an example of guest access topology using the H-REAP AP.

Figure 7-6 Branch Guest Access using H-REAP Central Switching

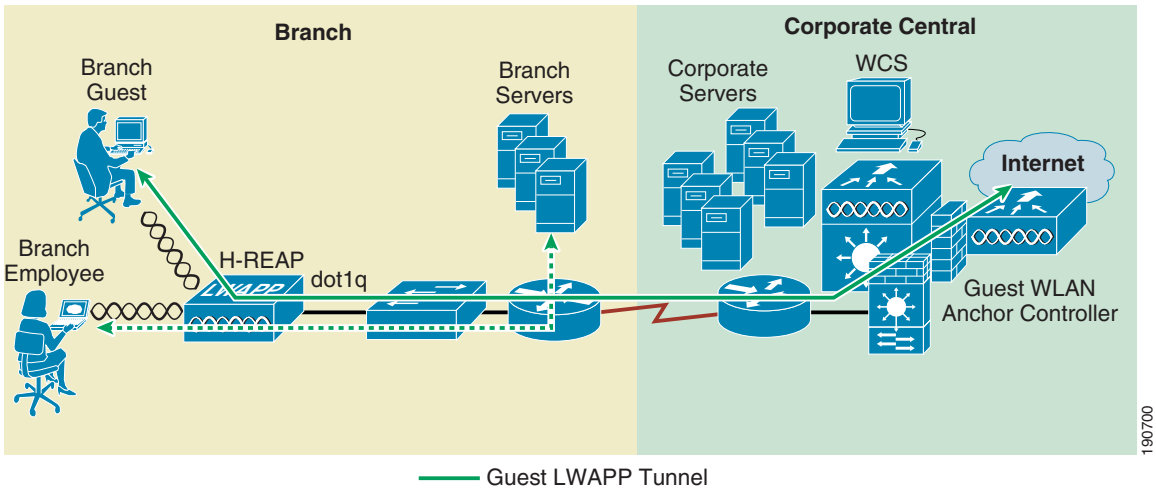

Public WLAN Hotspot

Many Public hotspot service providers are beginning to implement multiple SSID/WLANs. One reason for this is because an operator might want to offer an open authentication WLAN for web-based access and another WLAN that uses 802.1x/EAP for more secure public access.

The H-REAP AP, with its ability to map WLANs to separate VLANs, is now a alternative to an autonomous AP in small venue hotspot deployments where only one, or possibly two, APs are needed.

Figure 7-7 provides an example of hotspot topology using an H-REAP AP.

Figure 7-7 Hotspot Access using H-REAP Local Switching

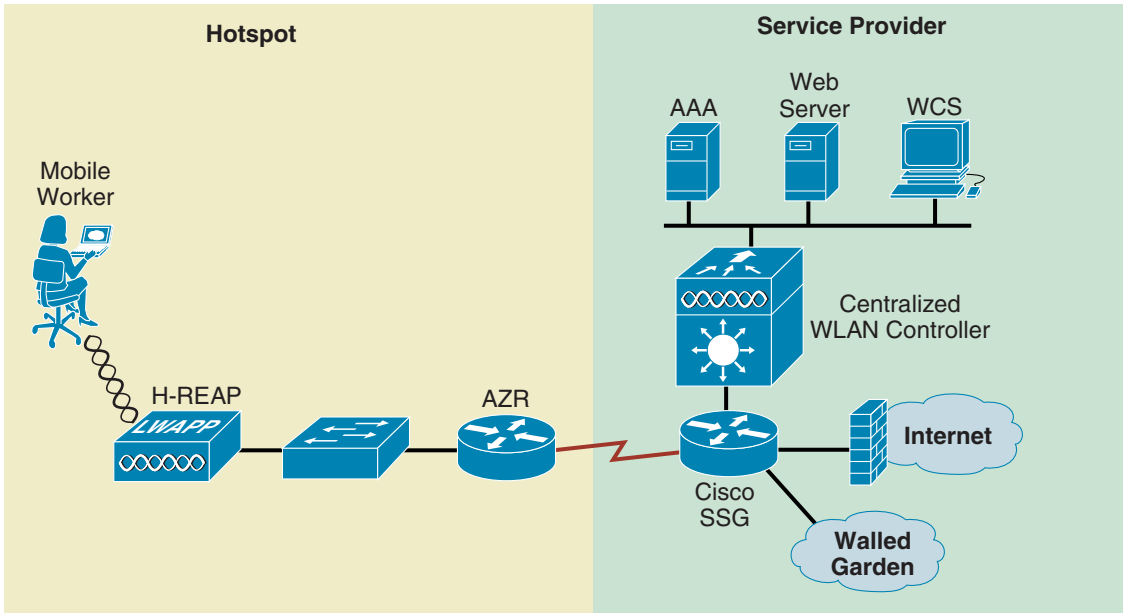

Deployment Considerations

The following section covers the various implementation and operational caveats that are associated with deploying H-REAP APs.

WAN Link

For the H-REAP AP to function predictably, there are a couple of things to keep in mind with respect to WAN link characteristics:

- Latency—A given WAN link should not impose latencies greater than 100 ms. The AP sends heartbeat messages to the controller once every thirty seconds. If a heartbeat response is missed, the AP will send five successive heartbeats (one per second) to determine if connectivity still exists. If connectivity is lost, then the H-REAP AP switches to standalone mode (see Operation Modes, page 7-3 for operation mode definitions). The AP itself is fairly delay tolerant. It is on the client where timers associated with authentication are sensitive to link delay and thus a constraint of ≤ 100 ms is required. Otherwise, the client could timeout waiting to authenticate, which, in turn, could cause other unpredictable behaviors, such as looping.
- Path MTU—WLAN controller software images 4.0 and later, applying to both the 1030 REAP and H-REAP APs, required an MTU no smaller than 500 bytes.

Authentication Methods

See Table 7-1 for a matrix of supported authentication methods based on the H-REAP mode of operation.

Table 7-1 Supported Authentication Modes

Authentication Method	Connected Mode	Standalone Mode	Notes
Open	Yes	Yes	
Shared	Yes	Yes	
EAP (TLS, PEAP, SIM and so on)	Yes	No	If the AP transitions to standalone mode, existing authenticated client sessions remain connected but no new authentications are possible.
WPA-802.1x	Yes	No	If the AP transitions to standalone mode, existing authenticated client sessions remain connected but no new authentications are possible.
WPA-PSK	Yes	Yes	
WPA2-802.1x	Yes	No	If the AP transitions to standalone mode, existing authenticated client sessions remain connected but no new authentications are possible.
WPA2-PSK	Yes	Yes	
Guest Access (Web Auth)	Yes	No	
VPN	Yes	No	
L2TP	Yes	No	
NAC	Yes	No	

Roaming

When an H-REAP AP is in Connected mode, all client probes (probe requests are handled at the AP, but are also forwarded to the WLC), association requests and response messages are passed between the H-REAP AP and the controller via the LWAPP control plane. This is also true for open, static wep, and wpa psk-based WLANs even though LWAPP connectivity is not required to support those authentication methods.

- Dynamic WEP /WPA—A client that roams between H-REAP APs using one of these key management methods must perform full authentication each time it roams, except in cases where the client supplicant supports Cisco CCKM. Otherwise, full 802.1x authentication is required (based on some EAP method) via the LWAPP control plane to an upstream AAA. After successful authentication, new keys are passed back to the AP and client. This behavior is the same as that in a traditional centralized WLAN deployment, except that in an H-REAP topology there can be link delay variations across the WAN, which can in turn impact total roam time.
- WPA2—To improve client roam times, WPA2 introduced key caching capabilities, based on IEEE's 802.11i specification. Cisco created an extension to this specification called Proactive Key Caching, or PKC. PKC today is supported only by Microsoft's Zero Config Wireless supplicant and Funk's (Juniper) Odyssey client. Cisco CCKM is also compatible with WPA2.

  LWAPP APs support PKC—PKC capable clients that roam between LWAPP APs do not perform full 802.1x authentication. Instead, the client and AP recompute their PMKID using the PKC method and immediately begin key exchange. It should be noted that this exchange occurs between the AP and the controller via the LWAPP control plane so, while roam times can be improved, there is still a potential for link delay variations across the WAN, and it should also be noted that PKC is not supported for locally switched WLANs. Remote branch locations requiring predictable, fast roaming behavior should consider deploying a local WLAN controller, such as the Cisco WLC2006 or NM-WLC for Integrated Service routers.
- Cisco Centralized Key Management (CCKM)—H-REAP APs currently do not support CCKM fast roaming. As such, CCKM clients will undergo full 802.1x authentication every time they roam from one H-REAP to another.
- Layer 2 Switch CAM Table Updates—When a client roams from one AP to another on a locally switched WLAN, the H-REAP AP currently does not announce to a Layer 2 switch that the client has changed ports. The switch does not discover that the client has roamed until the client performs an ARP for its gateway router. This behavior, while subtle, can have an impact on roaming performance.

Note H-REAP clients roaming on local switched WLANs where the APs reside on different subnets must renew their IP addresses when roaming to ensure they have an appropriate address for the network to which they have roamed.

WAN Link Disruptions

As described in sections Operation Modes, page 7-3 through H-REAP States, page 7-4, certain H-REAP modes and functionality require LWAPP control plane connectivity to the controller. Following is a summary of the features and functions that are impacted when the H-REAP is in Standalone mode.

EAP 802.1x and Web Auth WLANs

If existing local switched clients remain connected until the client roams or session re-authentication. No new client authentications are permitted.

If existing central switched clients are disconnected, no new client authentications are permitted.

As mentioned in H-REAP States, page 7-4, open, static WEP, and WPA/2-PSK configured WLANs can function in either Connected or Standalone modes and therefore are not impacted in the same way as WLANs requiring RADIUS services, such 802.1x or web authentication. If there is a requirement for a remote branch location to maintain wireless connectivity during WAN link disruptions, we suggest that a backup WLAN be implemented based on one of the three Layer 2 security polices above. Of these, WPA/2-PSK offers the strongest security and therefore is strongly recommended.

Other Features

The following features are unavailable when an H-REAP is in standalone mode:

- Radio resource management DFS support is maintained
- Wireless intrusion detection
- Location-based services
- NAC
- Rogue detection
- AAA override

Radio Configuration

The following radio configuration information is maintained when an H-REAP is in standalone mode:

- DTIM
- Beacon period
- Short preamble
- Power level
- Country code
- Channel number
- Blacklist

H-REAP Limitations and Caveats

Local Switching Restrictions

If one of the following VPN security methods is configured on the controller for a specific WLAN, then that WLAN cannot be configured for local switching for use by an H-REAP AP:

- IPSEC
- L2TP
- PPTP
- CRANITE
- FORTRESS[1]

Note VPN pass-through to external aggregation platforms is permitted. However, controller-imposed VPN passthrough restriction is not permitted.

Max Supported WLANs

H-REAP APs support eight WLANs. Therefore, any WLAN that is expected to be supported by an H-REAP AP must fall within WLAN IDs 1–8. WLAN IDs 9–16 are not propagated.

Network Address Translation (NAT/PAT)

Controller

A controller cannot reside behind a NAT boundary when communicating with APs because LWAPP APs communicate with the controller in two phases using two different IP addresses:

- Controller discovery—An LWAPP AP initially queries a list of controllers using a controller's management IP address. The management IPs are learned via DHCP Option 43, DNS, or they can be configured manually (see Initial Configuration, page 7-13). The discovery phase is used to determine which controller, within the list of eligible controllers, the AP will join. This is conveyed by sending an LWAPP control message containing the eligible controller's AP management IP address.
- Controller join—The AP joins the eligible controller using the learned AP management IP address. The AP management IP address cannot be supported by NAT because the AP learns this address during the discovery phase. Even if 1:1 NAT relationships are established, the controller is not capable of passing the AP manager's outside NAT address as the IP address the AP should use to join the controller.

AP

See Figure 7-8. A REAP or H-REAP AP can reside behind a NAT boundary in either of the following scenarios:

- If only one H-REAP AP resides behind a boundary, then PAT (NAT overload) can be used so long as port forwarding is enabled to map UDP ports 12222 and 12223 to the inside IP representing the H-REAP AP. We strongly recommend that the H-REAP be configured with a static IP address, or be given a static DHCP reservation to ensure that the NAT port mapping function works reliably. 1:1 NAT is also an option.
- If more than one H-REAP AP resides behind a boundary, only static 1:1 (inside to outside) NAT mapping can be used. Otherwise, PAT also operates correctly because of the unique UDP source port used by each AP. Multicast LWAPP messages do not correctly traverse NAT or PAT.

Figure 7-8 H-REAP with NAT/PAT

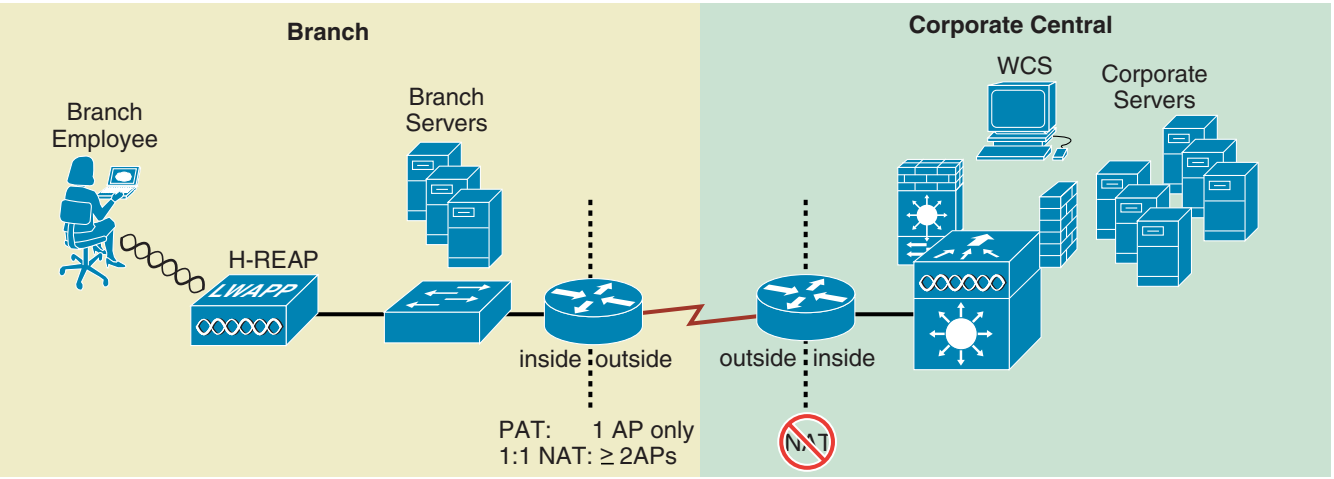

RADIUS Assigned VLANs

RADIUS-based VLAN assignment is supported for those H-REAP WLANs that are central-switched. This feature is not available when the H-REAP is in Standalone mode.

Web Authentication (Guest Access)

Controller-based web authentication may be used with local switched WLANs so long as the H-REAP is in Connected mode. Otherwise, those WLANs using web authentication are unavailable when the H-REAP is in Standalone mode.

Restricting Inter-Client Communication

Two or more clients, associated to a WLAN that is locally switched (by an H-REAP), are not prevented from communicating with one another even if Peer-to-Peer Blocking mode is enabled on the controller. This is because locally switched wireless traffic does not go through the controller. If it becomes necessary to block inter-client communication for a local-switched WLAN, then some kind of uRPF, such as ACL, can be applied at the ingress interface of the first Layer 3 hop.

Those H-REAP WLANs that are central switched have inter-client communication restricted based on the Peer-to-Peer Blocking mode setting on the controller.

H-REAP Scaling

- Per-Site—Sites requiring more than three APs should consider deploying a controller locally at the branch location. There are a few reasons for this:
  - Roaming performance—As described in Roaming, page 7-9, roaming performance can be impacted by the availability and link characteristics of the WAN backhaul. This is true even when key caching methods, such as 802.11i or Cisco CCKM, are employed.
  - Reliability—Branch WLAN topologies that depend on authentication, radio resource management and other upstream services are only as good as the availability of the WAN backhaul.

  - WAN backhaul bandwidth consumption—As the number of H-REAP APs increases, bandwidth use also increases as a result of LWAPP control plane traffic. All client probes, association requests, and authentication-related messages result in LWAPP control and data traffic being sent across the WAN to the controller, even when WLANs are local switched.
- Per-controller—There are no restrictions with regard to the number of APs that can operate in H-REAP mode. The total number of H-REAP APs per controller is bound only by the maximum number of Lightweight APs that are supported for a given controller model.

Inline Power

The Cisco 1130 and 1240 Series APs support both the Cisco inline power specification and conform to the 802.3af standard, whereas the former Cisco 1030 Series REAP APs support 802.3af only.

Management

H-REAP APs can be managed and monitored either through the controller's GUI or Cisco Wireless Control System (WCS) in the same way that regular LWAPP APs are managed. The only exception is when the H-REAP APs become un-reachable due to WAN outages. For more information on management and WCS please see chapter X, section Y of this document.

H-REAP Configuration

Initial Configuration

An eligible Cisco 1130 or 1240 series AP requires the following minimum information to join a controller so that it can be configured for H-REAP operation:

- An IP address
- A default gateway address
- Management interface IP address of one or more controllers

The above information can be obtained in one of four ways:

- Static configuration via serial console port
- DHCP with statically configured controller addresses
- DHCP with DNS resolution for controller addresses, as discussed in Chapter 2, “Cisco Unified Wireless Technology and Architecture”
- DHCP with Option 43, as discussed in Chapter 2, “Cisco Unified Wireless Technology and Architecture”

Serial Console Port

Unlike the earlier 1030 series REAPs, The 1130 and 1240 series APs offer a serial console port that can be used to establish basic parameters for connectivity. Use the following steps to establish initial configuration using the console port method. The serial console port method can be used only when the AP is not actively joined with a controller when being configured and is running LWAPP image 12.3(11)JX or later.

Note Complete Step 4 a. and b. only if DHCP will not be used at the branch to assign an IP address to the H-REAP AP. Care must be taken to ensure that the addresses used conform to the addressing scheme being used at the branch location.

Step 1 Using a standard Cisco DB9/RJ45 console cable connect the AP to a laptop running Hyper Terminal or other compatible terminal communications software. As with all Cisco devices, the serial parameters need to be set at 9600bps, 8 data bits, 1 stop bit and No flow control.

Step 2 Power on the AP. To configure the AP through the console port, it should not be connected to the network. Otherwise, if the AP discovers a controller and joins it, you will not be able to establish an exec session. Therefore, Cisco recommends that the AP remain disconnected from the network (Standalone mode) until the initial configuration has been completed.

Step 3 After the AP has completed loading its local image, establish an exec session by typing **enable** and then entering **cisco** for the enable password.

Step 4 At the <ap-mac-address># prompt, use the following commands to configure the IP, mask, gateway, hostname, and the primary controller:

a. **lwapp ap ip address** *ip-addr subnet-mask*

b. **lwapp ap ip default-gateway** *ip-addr*

c. **lwapp ap hostname** *ap-hostname* (optional)

d. **lwapp ap controller ip address** *ip-addr*

Note If DHCP services are used within the branch (see DHCP with Statically Configured Controller IPs, page 7-15) and you do not want to use DHCP Option 43 or DNS methods to issue controller management IP addresses, enter only the **lwapp ap controller ip address** *ip-addr* command from Step 4.

The preceding commands are saved directly to NVRAM.

Step 5 To review the static configuration, type the following command:

show lwapp ip config

Output similar to the following is displayed:

```
AP0014.1ced.494e# sho lwapp ip config
LWAPP Static IP Configuration
IP Address         10.20.104.50
IP netmask         255.255.255.0
Default Gateway    10.20.104.1
Primary Controller 10.20.30.41

AP0014.1ced.494e#
```

If an error has been made, repeat the commands listed in Step 4 to correct.

Step 6 To clear one or more static entries, use the following commands:

a. **clear lwapp ap ip address**

b. **clear lwapp ap ip default-gateway**

c. **clear lwapp ap controller ip address**

d. **clear lwapp ap hostname**

When you are connected to the branch network, the AP boots and sends discovery requests to each controller defined in Step 4 d. The AP then joins the least used controller.

DHCP with Statically Configured Controller IPs

This method uses DHCP to dynamically configure the AP with an IP address and default gateway. The DHCP service can be implemented locally or remotely using an external server or locally using DHCP services resident within IOS. The WLC management interface IP addresses can be manually configured using the APs console interface; this can either be done before shipping to the branch office or on site. See Serial Console Port, page 7-13. When connected to the branch network, the AP boots and sends discovery requests to each controller defined. The AP then joins the least used controller.

Configuring AP for H-REAP Operation

The following configuration tasks are accomplished using the wireless LAN controller GUI interface.

When an AP joins the controller for the first time it defaults to local AP mode. The AP must be set for H-REAP mode before local switching parameters can be established.

Step 1 From the controller **Wireless** configuration tab, locate the newly joined AP and click **Detail** (see Figure 7-9:

Figure 7-9 Wireless Configuration Tab

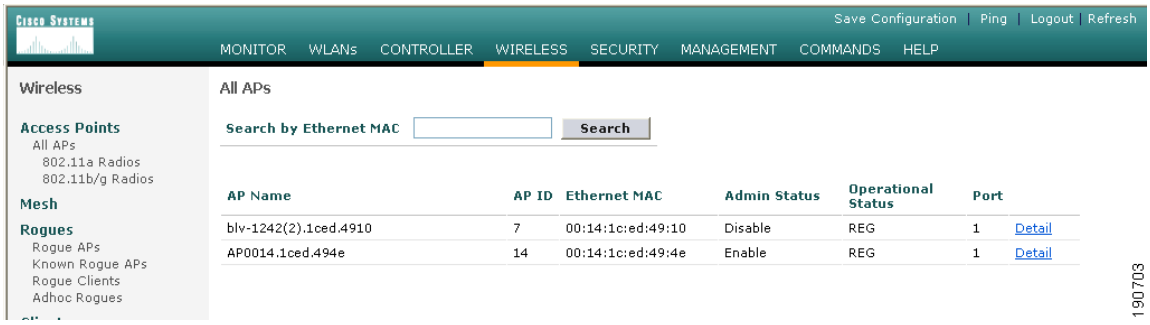

Step 2 Select **AP Mode**, **Name**, **Location**, and **Controller Priority**.

From the AP mode drop-down list, choose **H-REAP**. (See Figure 7-10.)

Figure 7-10 Wireless Configuration—AP Mode

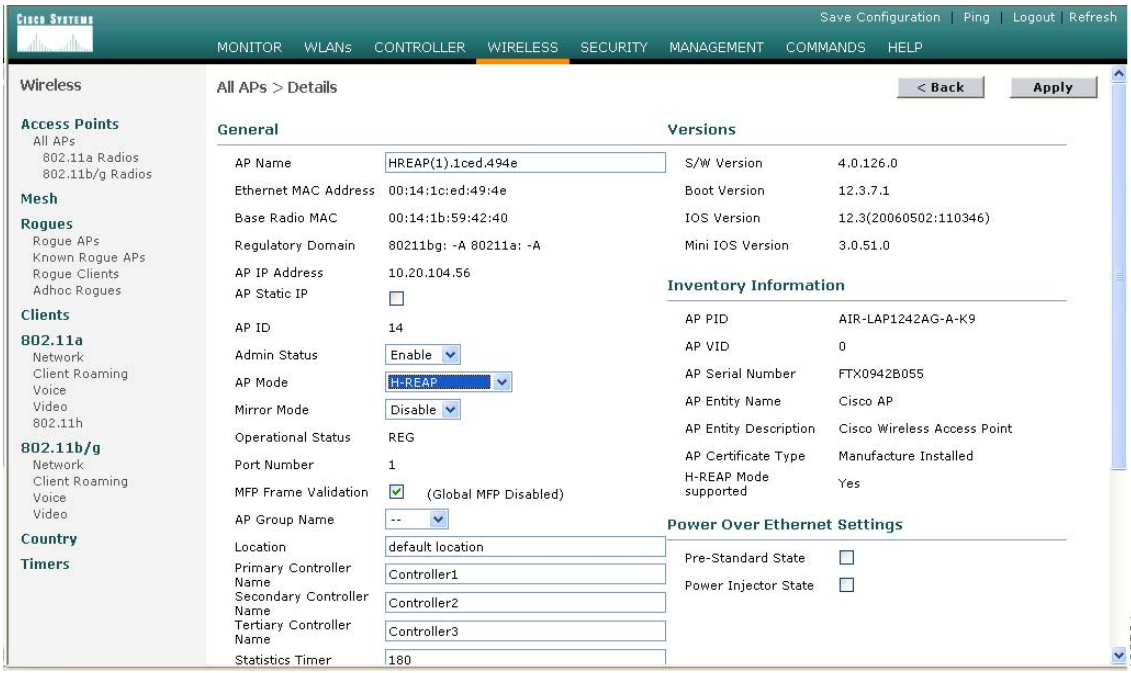

Optionally, configure an AP name or optionally configure a location name.

Step 3 Identify the primary controller the AP should join and, optionally, a secondary and tertiary controller in the event the primary (or secondary) controller becomes unreachable.

These names are case-sensitive and correspond to the system name. If none of the named controllers are available, the AP will join one of the other controllers that belong to the mobility group based on automatic load balancing.

Step 4 Click **Apply**.

The AP reboots and re-joins the controller in H-REAP mode.

Note When the H-REAP AP reboots, its interface is not configured for 802.1q trunking mode. Ensure that the DHCP scope used for assigning addresses to H-REAP APs is configured on the native VLAN because the AP originates DHCP requests with no VLAN tag.

Enabling VLAN Support

After the H-REAP AP has joined the controller in H-REAP mode:

Step 1 Find the AP under the controller Wireless settings and click **Details**.

Note that there are new H-REAP configuration settings presented in the AP details window. (See Figure 7-11.)

Step 2 Place a check mark in the **VLAN Support** check box.

Note that a Native VLAN ID definition window and a VLAN Mappings button are added.

Step 3 Enter the VLAN number defined as the native VLAN.

Step 4 Click **Apply**.

Figure 7-11 Wireless Settings

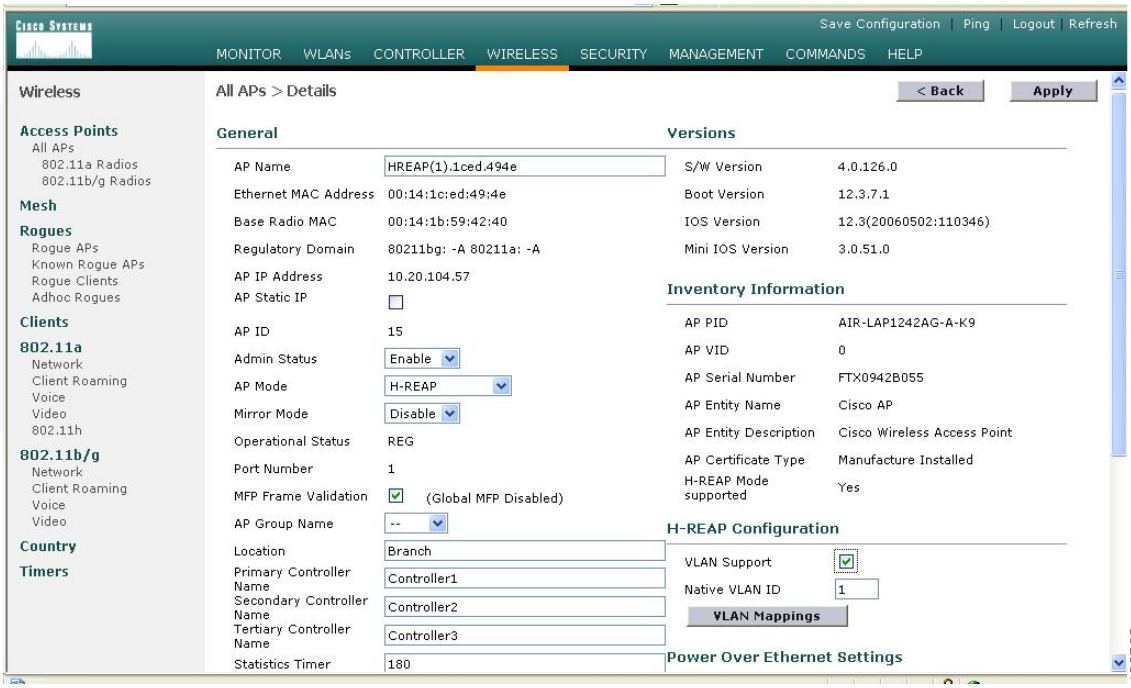

Advanced Configuration

The following steps outline how to configure an H-REAP AP to perform local and or central switching in addition to highlighting any caveats associated with the configuration process.

Choosing WLANs for Local Switching

Before a WLAN can be mapped to a local VLAN on the H-REAP AP, the WLAN must first be made eligible for H-REAP local switching.

Step 1 Click the **WLANs** tab.

Step 2 Find the WLANs that need to be locally switched and click **Edit**. (See Figure 7-12.)

Figure 7-12 WLANs Tab

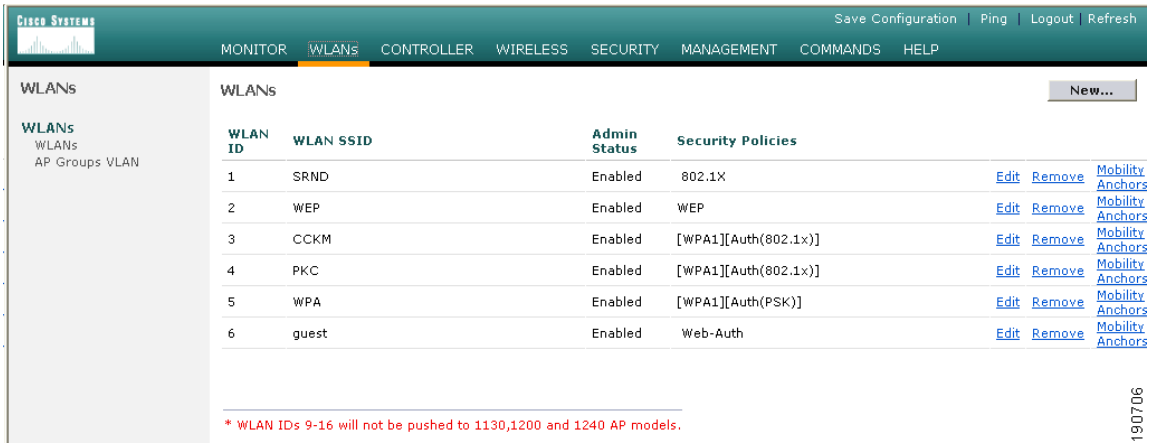

Configuring H-REAP Support on a WLAN

Step 3 Place a check mark in the **H-REAP Local Switching** check box. (See Figure 7-13.)

Figure 7-13 WLANs—Edit

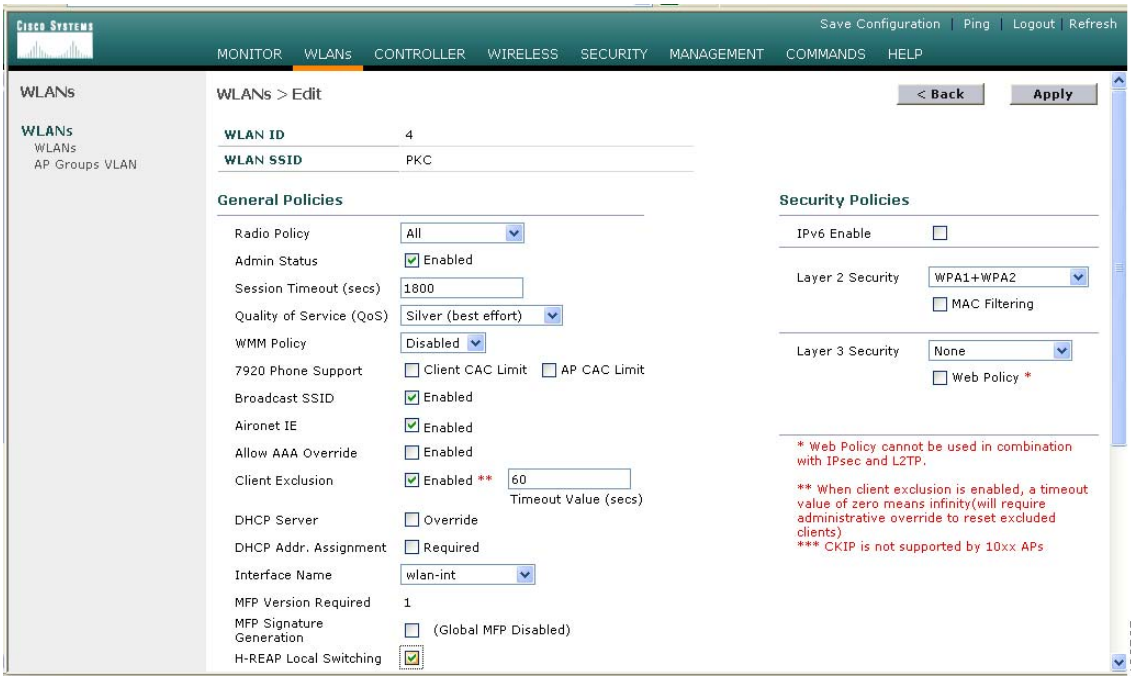

Step 4 Click **Apply**.

H-REAP Local Switching (VLAN) Configuration

After all eligible WLANs have been configured to support H-REAP, do the following:

Step 1 Click the **Wireless** tab.

Step 2 From the list of APs, find the H-REAP and click **Detail**. (See Figure 7-14.)

Figure 7-14 Wireless Tab—APs

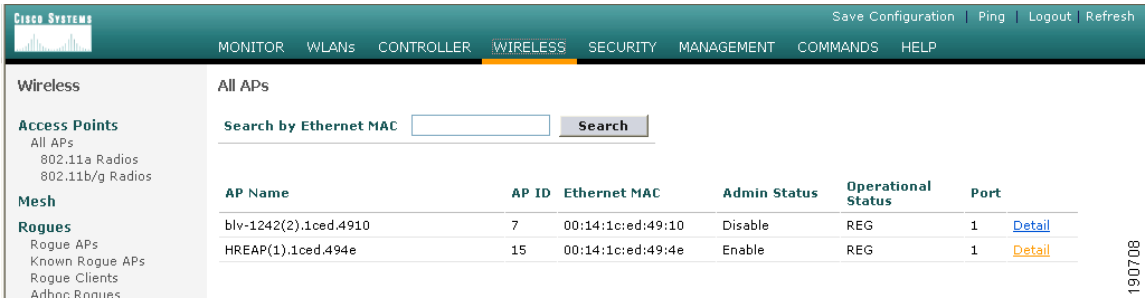

Step 3 From the AP Details configuration page, click **VLAN Mappings**. (See Figure 7-15.)

Figure 7-15 All APs—Details

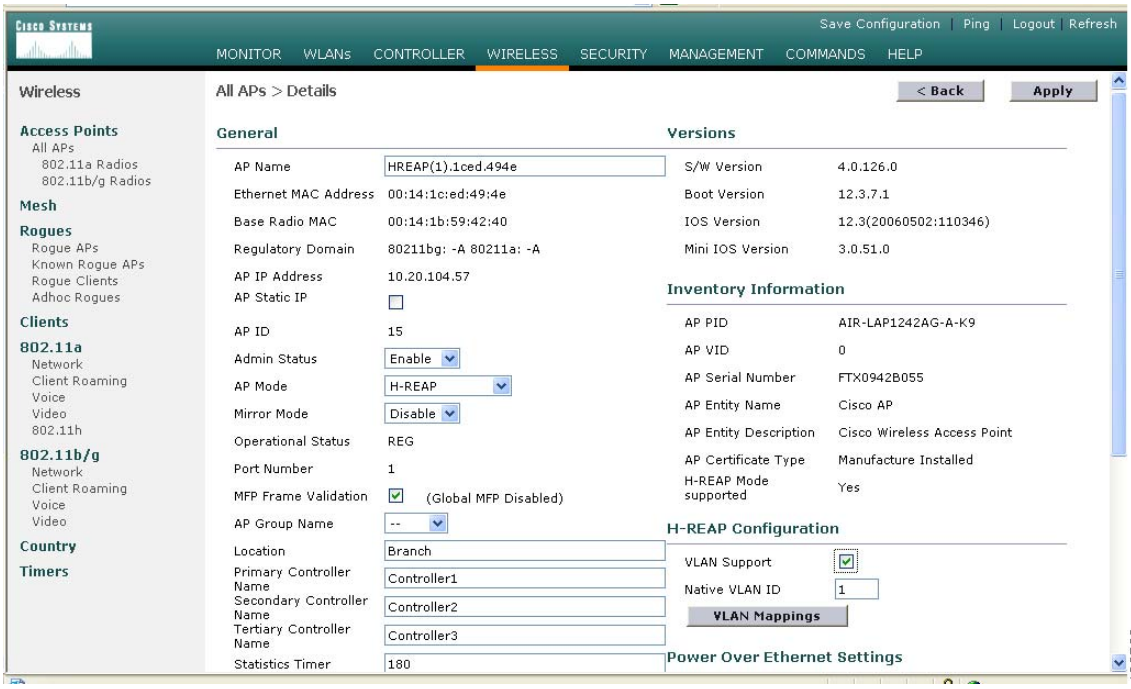

Establishing a WLAN to VLAN Mapping

The VLAN Mappings page displays all WLANs that have been configured for H-REAP switching, along with a configurable VLAN ID field. (See Figure 7-16.)

Figure 7-16 VLAN Mappings

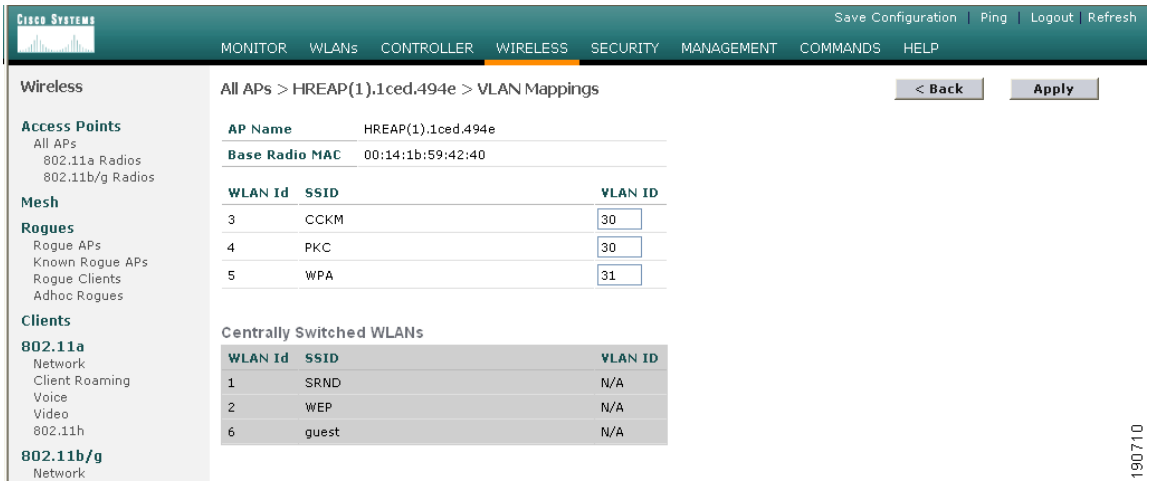

Note The WLAN IDs that are displayed initially are inherited from the central controller WLAN interface settings.

Step 1 For each WLAN/SSID, configure a locally relevant VLAN number.

More than one WLAN can be mapped to local VLAN number.

Step 2 Click **Apply**.

Note All WLANs shown in the grey box are centrally switched and might not be active, depending on whether the WLAN is enabled globally. All user traffic associated with a centrally switched WLAN are tunneled back to the controller.

Centrally switched WLANs can be excluded from the H-REAP by using the WLAN override feature to uncheck the WLANs that are not required.

Note For each locally switched WLAN, there must be an DHCP helper address or local DHCP pool configured for the mapped VLAN.

H-REAP Verification

Verifying the H-REAP AP Addressing

- If using DHCP to assign an address, verify DHCP server configuration settings, correct subnet, mask, and default gateway.
- Ensure AP DHCP scope is defined on the native VLAN.

- If AP was configured with a static addresses, ensure AP address, subnet, mask and gateway are consistent with addressing scheme used within the branch location using the **show lwapp ip config** command. See Serial Console Port, page 7-13 for more information.

Verifying the Controller Resolution Configuration

- If using DHCP Option 43/60 for controller resolution, verify that the VCI and VSA string format on the DHCP server is correct.
- Verify that the correct controller management IP address is configured in the DHCP server.
- If using DNS resolution, verify that a DNS query of CISCO-LWAPP-CONTROLLER@localdomain can be made from the branch location and resolves to one or more valid controller management IP address.
- Verify valid DNS server addresses are being assigned via DHCP
- If the controller IP was configured manually, verify the configuration via the serial console port with the AP disconnected from the network using the **show lwapp ip config** command. See Serial Console Port, page 7-13 for more information.

Troubleshooting

This section provides troubleshooting guidelines for some common problems.

H-REAP Does Not Join the Controller

If an H-REAP AP is not joining the expected controller:

- Verify routing from the branch location to the centralized controller. Check that you can ping the Controller management IP address from the AP subnet.
- Verify that the LWAPP protocol (UDP ports 12222 and 12223) is not being blocked by an ACL or firewall
- Verify that the H-REAP hasn't joined another controller in the mobility group

  Check to see whether a controller within the mobility group has been designated as “master controller”, which could cause an H-REAP to join a controller other than the one expected.

Client Associated to Local Switched WLAN Cannot Obtain an IP Address

- Verify that 802.1q trunking is enabled (and matches the AP configuration) on the switch and or router ports to which the AP is connected.
- Verify that an IP helper address or local DHCP pool has been configured for the VLAN (sub-interface) at the first Layer 3 hop for the WLAN in question.

Client Cannot Authenticate or Associate to Locally Switched WLAN

If local switched WLAN uses central authentication:

- Verify H-REAP is not in Standalone mode (WAN backhaul down).
- Verify a valid RADIUS authentication server has been configured for the WLAN.
- Verify reachability to the RADIUS authentication server from the controller.

- Verify that the RADIUS server is operational.
- Verify that the authentication service and user credentials are configured on the RADIUS server.

If the local switched WLAN uses a pre-shared key:

- Verify that the WPA or WEP configuration on the client matches that of the WLAN.
- Verify if wireless client requires WLAN SSID to be broadcast (if disabled) to authenticate/associate.

Client Cannot Authenticate or Associate to the Central Switched WLAN

If the central switch WLAN uses central authentication:

- Verify H-REAP is not in Standalone mode (WAN backhaul down)
- Verify a valid RADIUS authentication server has been configured for WLAN
- Verify reachability to RADIUS authentication server from the Controller
- Verify that the RADIUS server is operational.
- For AAA authenticated clients, verify that authentication service and user credentials are configured on the RADIUS server.

If local switched WLAN uses a pre-shared key:

- Verify that the WPA or WEP configuration on the client matches that of the WLAN.
- Verify if wireless client requires WLAN SSID to be broadcast (if disabled) to authenticate / associate.

H-REAP Debug Commands

This section contains debug commands that can be used for advanced troubleshooting.

Controller Debug Commands

The following commands are entered through, and their output can be viewed using, the controller's serial console interface:

debug lwapp events enable

debug lwapp packets enable

H-REAP AP Debug Commands

The following commands are entered through, and their output can viewed using, the H-REAP serial console interface:

debug lwapp client packet

debug lwapp client mgmt

debug lwapp client config

debug lwapp client event

debug lwapp reap load

debug lwapp reap mgmt

Cisco Unified Wireless Control System

Introduction

The modern day Wi-Fi 802.11 wireless network has evolved to become an integral part of the overall enterprise infrastructure, and because such organizations are seeking similar capabilities from their wireless systems management as they have from their enterprise infrastructure management systems in the past. IT managers and other networking professionals expect capabilities in such tools, enabling them to ensure their mission-critical wireless network systems are reliable, available, and performing optimally. A robust and reliable centralized network management solution capable of uniformly managing geographically disparate WLANs is necessary to simplify operations and to reduce total cost of ownership.

This chapter describes the Cisco Wireless Control System (WCS) and addresses management considerations that you should consider when using it to design, deploy, and manage your enterprise wireless LAN. It is intended for the reader responsible for performing such tasks using Cisco Unified Wireless Network (UWN) technology.

The following sections discuss various areas of WLAN management including wireless LAN configuration and monitoring, RF management and system planning, intrusion monitoring, and location tracking as follows:

- Wireless Control System Overview—Describes network management in general, along with a brief overview of the Cisco Wireless Control System (WCS).
- Role of WCS Within the Unified Wireless Network Architecture—Describes the overall network architecture and illustrates management data flows.
- How WCS can be used to define and configure devices within your wireless network.
- Using WCS to Monitor Your Wireless Network—Discusses how to use WCS to monitor your network in daily operation. A detailed explanation of the relationship between traps, events, alarms, and notifications can be found in this section and should be valuable to anyone considering using WCS to alert management and other personnel.
- Using WCS to Locate Devices in Your Wireless Network—Examines how WCS can provide on-demand location of WLAN clients, asset tags, and rogues. The use of the Cisco Wireless Location Appliance is also discussed with references provided to comprehensive sources of information on Cisco Location-Based Services (LBS).
- Using WCS to Efficiently Deploy Your Wireless Network—Suggests aspects of WCS that you can use to assist you with efficient deployment of a multi-site wireless network.

- Traffic Considerations When Using WCS in Large Networks—Discusses the traffic generated by polling and other sources within WCS. Those users planning very large, multi-server implementations over remote networks should consider the information in this section when making planning and design decisions.
- Administering WCS—Examines WCS scheduled tasks, users, and database administration.

Wireless Control System Overview

The Cisco Wireless Control System (WCS) is a component of the Cisco Unified Wireless Network (UWN) that provides a powerful network management solution allowing the design, control, and monitoring of enterprise wireless networks from a centralized location. The benefit of this is simplified operations and reduced total cost of ownership because network administrators now have a single solution for RF prediction, policy provisioning, network optimization, troubleshooting, user tracking, security monitoring, and general WLAN systems management. WCS enhances the management and control capabilities already present in the Cisco UWM via the WLAN controller web user interface and command line interface (CLI).

WCS makes it possible for the point of control in an enterprise WLAN to move from individual controllers to a network of controllers. WCS provides graphical views of multiple controller hardware formats and offers a comparable level of configuration, performance monitoring, accounting, security, and fault management to that offered at the controller level.

WCS functionality can be grouped into the following areas:

- Network monitoring and troubleshooting

  Cisco WCS provides tools that enables the visualization of wireless networks as well as the monitoring of ongoing WLAN performance. Cisco WCS also provides a portal into the real-time RF management capabilities provided by Cisco wireless LAN controllers including automated channel assignments and access point transmit power settings. Quick visibility into coverage holes, device status alarms, and key usage statistics is provided for easy WLAN monitoring and troubleshooting.

- Indoor location tracking

  Cisco WCS provides you with the ability to efficiently track wireless devices, including Wi-Fi enabled laptops, PDAs, and voice handsets as well as mobile assets equipped with Wi-Fi 802.11 active RFID tags. The base version of WCS can determine with which access point a wireless device is associated and provides a general idea of where wireless devices are situated. Environments that require more granular location services can optionally license additional location-based services capabilities within WCS and take advantage of Cisco RF Fingerprinting technology, which is capable of providing accuracy to 10 meters or better. To scale the use of location tracking beyond single threaded device localization, Cisco WCS with location can be deployed in conjunction with the Cisco Wireless Location Appliance for real-time simultaneous tracking of up to 2500 wireless devices.

- Wireless LAN planning and design

  Integrated RF prediction tools are available to create detailed wireless LAN designs, including lightweight access point placement, configuration, and performance/coverage estimates. Floor plans can be imported into Cisco WCS and RF characteristics assigned to building components to increase design accuracy. Graphical heat maps help visualize anticipated wireless LAN behavior to facilitate planning and deployment.

- Policy management and enforcement

  A full suite of tools is provided for the management and enforcement of security policies within a Cisco wireless infrastructure, including the following:

  - *Support for the Cisco Unified Intrusion Detection System/Intrusion Prevention System*—When used with a Cisco Unified IDS/IPS (part of the Cisco Self-Defending Network), the IDS/IPS device detects when an associated client sends malicious traffic through the Cisco Unified Wireless Network and sends shun requests to Cisco Wireless LAN Controllers. These controllers then in turn disassociate the client device.
  - *RF attack signatures and wireless intrusion prevention*—Customizable attack signature files can be used to rapidly detect common RF-related attacks such as denial of service (DoS), Netstumbler, and FakeAP. Cisco WCS is capable of raising alarms and generating notifications if an attack is detected. Detailed trending reports enable network administrators to identify recurring security threats before they can cause significant harm to the network.
  - *Rogue detection, location, and containment*—Cisco WCS maintains a constant vigil for unauthorized "rogue" access points, ad-hoc networks, and clients. If unauthorized rogue devices appear, Cisco WCS can be used to determine their location and assess the level of threat. If deemed malicious, containment procedures can be initiated by the WCS operator to limit the potential threat posed by these devices.
  - *Policy creation and enforcement*—Cisco WCS contains a service policy engine that allows network administrators to easily create and enforce a wide variety of network policies including virtual LAN (VLAN), RF, quality of service (QoS), and security policies. Multiple WLANs can be created with unique service set identifiers (SSIDs) and individualized security parameters. These security policies can be applied across an entire Unified Wireless Network, to specific wireless LAN controllers, or even to individual lightweight access points.
  - *User exclusion lists*—Cisco WCS can be used to proactively exclude specific users from associating with the wireless network. If unusual activity is detected, offending devices can quickly be flagged and excluded if considered to be malicious. These devices cannot access wireless LAN services until a pre-configured timer has expired or a manual override is initiated to grant wireless LAN access once again.

- Secure guest access

  Cisco WCS allows customizable guest access capabilities that allow organizations to keep their wireless networks secure while providing customers, vendors, and partners with controlled access to their WLANs. Organizations can enable the *Guest Access Lobby Ambassador* feature on the wireless LAN controller to allow for the creation of local usernames and passwords and for local or RADIUS-based authentication of guest users.

- General wireless LAN systems management
  - *Configuration*—Configuration of network components can be done via traditional manual methods on an individual basis, or WCS administrators can assign a template to one or all of the wireless LAN controllers or access points in a mobility group.
  - *Troubleshooting*—Important network information is consolidated and reported on, such as noise levels, signal-noise ratio, interference, and signal strength. This facilitates isolation and resolution of problems at all layers of a wireless network.
  - *Software updates*—Upgrades to software contained on components within the Cisco Unified Wireless Network can be performed from a centralized location.
  - *Customized reports*—Numerous reports are available that document network and system activity. These include client statistics, radio usage data, 802.11 counters, RF management configuration history, and device alarms.

Role of WCS Within the Unified Wireless Network Architecture

The Cisco Unified Wireless Network is designed to provide robust 802.11 wireless networking solutions for large enterprises, branch, and remote offices as well as outdoor areas. The system manages all data client communications and system administration functions, performs radio resource management (RRM), and manages system-wide mobility policies.

In this solution, the various Cisco WLAN controllers (embedded and standalone) together with their registered lightweight access points may be managed via the following:

- A controller web and command line interface (CLI)
- The Cisco Wireless Control System (WCS), which can be used to configure and monitor one or more controllers and registered access points. All Cisco wireless LAN controller models can be managed by Cisco WCS including enterprise-class standalone wireless LAN controllers such as the 4400 and 2000 Series; as well as the Cisco Catalyst 6500 Series Wireless Services Module (WiSM), the Cisco Catalyst 3750G Integrated Wireless LAN Controller, and the Cisco Wireless LAN Controller Module (WLCM) for Integrated Services Routers (ISRs)
- Other management software compliant with industry-standard SNMP v1, v2c, and v3 interfaces

Figure 8-1 shows the interoperation of WCS along with the other components of the Cisco Wireless LAN Solution when deployed in a Cisco Unified Wireless Network.

Figure 8-1 Overall Wireless Network Architecture with WCS

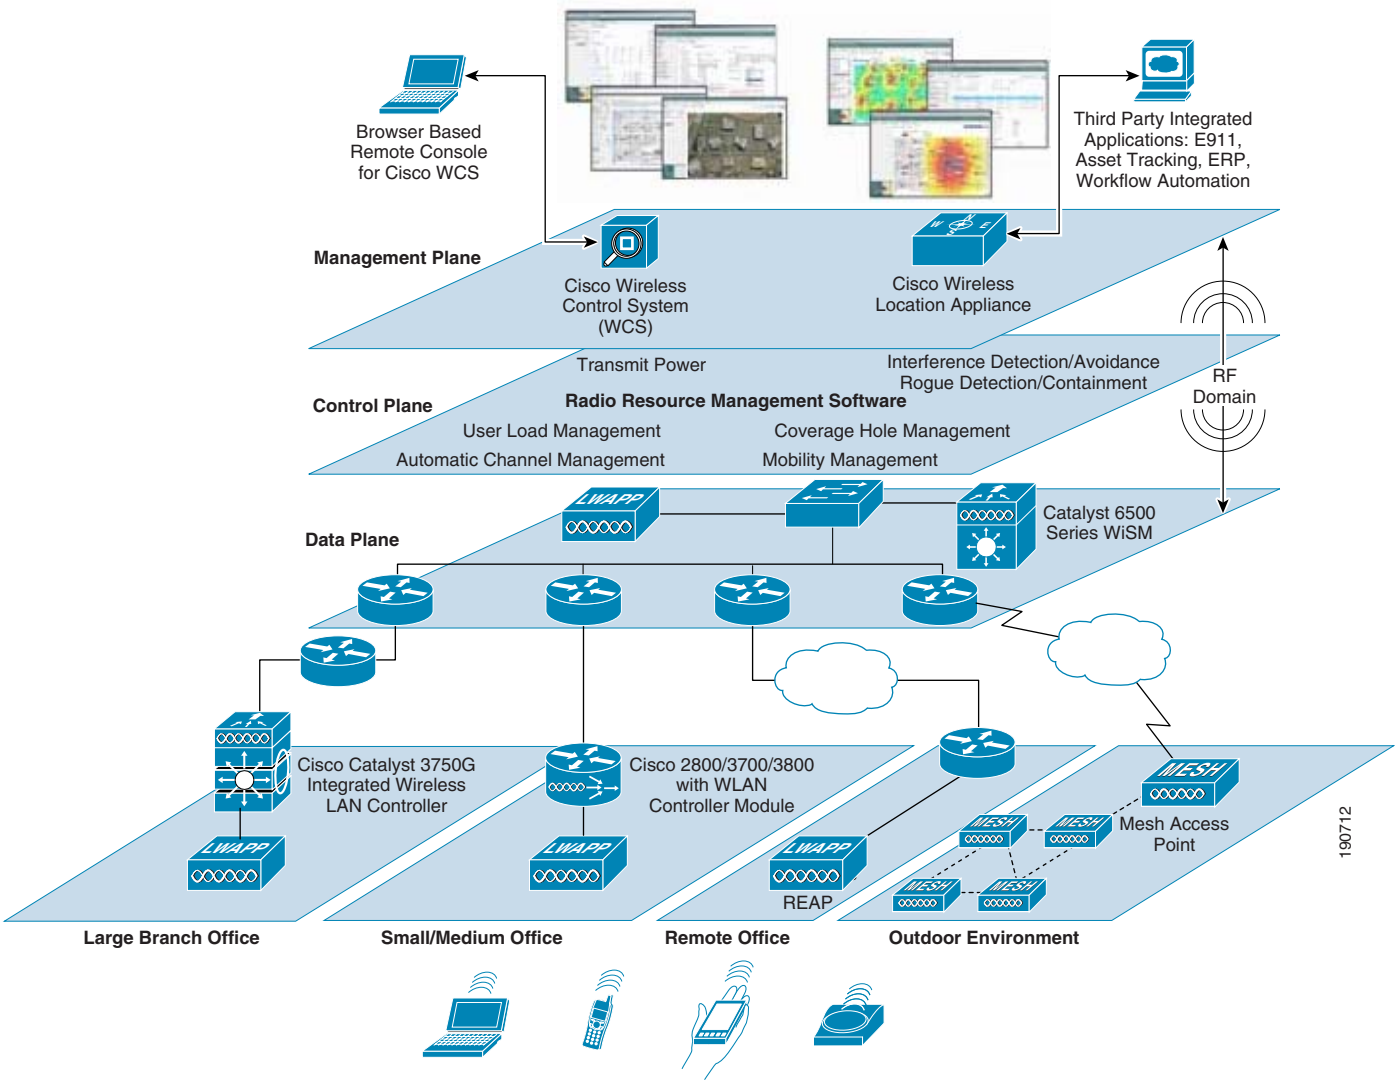

Various communication protocols (SNMP, SMTP, HTTP/HTTPS, FTP, TFTP, SOAP/XML, and so on) are implemented between these components to provide the management, alerting, and notification functionality necessary to efficiently manage modern enterprise wireless infrastructures.

Figure 8-2 shows the typical client/server communication flows between WCS, the client workstation browser, and the infrastructure components comprising the enterprise wireless LAN. WCS does not directly manage lightweight access points but rather communicates with SNMP agents contained within the wireless LAN controllers to which lightweight access points have been assigned. Configuration changes, inquiries, monitoring, and reporting are all handled via the exchange of SNMP traps, commands, and responses between WCS and the WLC SNMP agents. Any information or configuration requests concerning controller or access point resources are sent by WCS to these WLC SNMP agents. Working in conjunction with other controller hardware and software, these SNMP agents participate in the initiation of appropriate actions on internal controller resources, or communicate such actions to assigned lightweight access points via the lightweight access point protocol (LWAPP).

Figure 8-2 Management Data Flows within the Cisco Unified Wireless Network

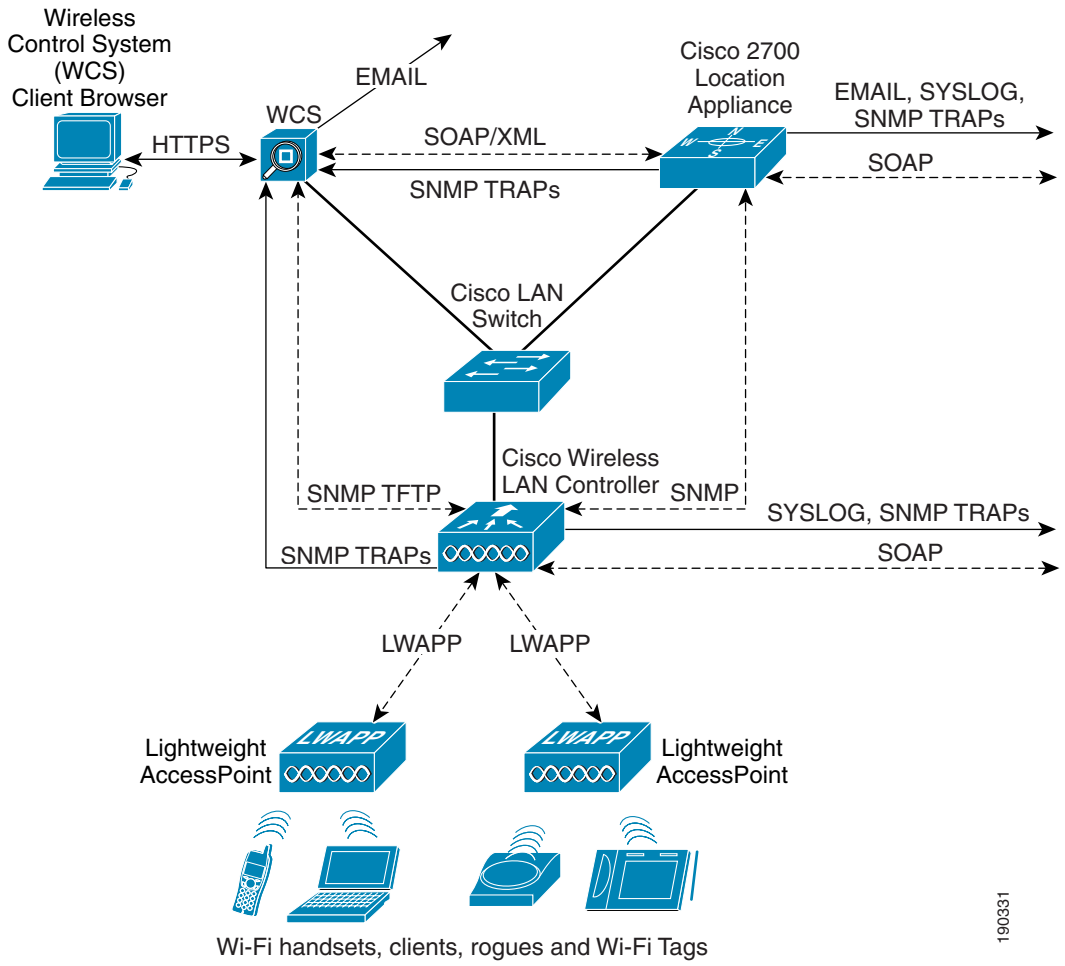

Figure 8-2 also shows the ability of controllers to transmit SNMP *traps* to up to six *trap receivers* as well as transmitting *syslog messages* to a remote syslog receiver. The ability to send traps to multiple trap receivers is useful in networks that, in addition to WCS, possess an overall enterprise network management system (NMS) that you would like to inform when traps are generated by the wireless network devices. The multiple trap capability in WLCs allow you to send traps to WCS and the enterprise NMS.

The information contained within SNMP traps and polling responses are the foundation of *events,* and based on their severity, these events can result in the triggering of WCS *alarms*. Depending on the configuration of WCS alarm *notification*, WCS can generate messages to e-mail destinations such as desktop and laptop clients, pagers, PDAs, and other systems notifying them of newly-triggered critical and coverage hole alarms.

File transfer protocols are used to update a variety of WLC software and configuration information. WCS readily accommodates this by providing for integrated TFTP and FTP server capability. WCS can also be used to configure TFTP file transfer between wireless LAN controllers and other TFTP servers that may be located closer on the network to the managed devices.

Real-time RF management is a hallmark feature of the Cisco lightweight wireless solution, and a unique product differentiator. Each WLAN controller uses dynamic algorithms to create an environment that is completely self-configuring, self-optimizing, and self-healing, making a Cisco-powered WLAN ideal for the delivery of secure and reliable business applications. This is done via specific *radio resource management (*RRM*)* functions such as the following:

- Radio resource monitoring
- Dynamic channel assignment
- Interference detection and avoidance
- Dynamic transmit power control
- Coverage hole detection and correction
- Client and network load balancing

Note Further information about RRM can be found in the RF Design chapter of this SRND. Additional information on this topic can also be found in various documents available at the following URL: http://www.cisco.com.

The Cisco WCS allows for straightforward configuration of RRM parameters that can be applied to multiple WLAN controllers using the policy template facility. If controllers detect that one or more various predefined RRM thresholds are violated, a trap is sent to the Cisco Wireless Control System (WCS). WCS provides multiple reporting facilities that can be used to view the RF environment in real time, aiding greatly in understanding what is happening in the air space and facilitating the troubleshooting process.

WCS provides both the user and control interfaces to the Cisco Wireless LAN Location Appliance, allowing simultaneous location display of WLAN clients, asset tags, rogue access points, and rogue clients. Additionally, WCS provides the ability to configure the location appliance to send various forms of user notification when changes occur in client or asset location. In this way, the location appliance can be defined to transmit messages using SOAP, SMTP, SNMP traps, or syslog messaging if clients or assets become missing, enter or leave coverage areas, or stray beyond a set distance from a pre-determined marker.

For complete information about WCS hardware and software requirements, and for complete step-by-step guidance on installing and accessing the WCS server, see the following documents:

- *Cisco Wireless Control System Configuration Guide, Release 4.0*— http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_book09186a00806b57ec.html
- *Cisco Wireless Control System Release Notes, Release 4.0*— http://www.cisco.com/en/US/products/ps6305/prod_release_note09186a00806b0811.html

Defining Network Devices to WCS

Before being able to manage WLAN controllers and location appliances, these devices must be defined to WCS. You need to specify the IP and SNMP information necessary to communicate with each device that you wish to include in the management domain of your WCS server. After these devices are defined, they are considered to be within the *management domain* of that WCS. WCS implicitly learns of the existence of any lightweight access points registered to any WLAN controllers defined to it.

When using commands in the following subsections that allow multiple objects to be selected as targets, the selected objects must all be present on one display page. This is important, for example, if the total population of controllers displayed spans several pages and requires paging forward and backward. The selected objects cannot be present across multiple pages.

Adding Controllers to WCS

Adding Controllers

Before being defined to WCS, all WLAN controllers should be properly configured as per the *Cisco Wireless Control System Configuration Guide, Release 4.0*. Basic communication settings for each deployed WLAN controller such as IP addressing, SNMP communities, strings and passwords, SNMP version in use, and so on, should be noted before attempting to define these resource to WCS.

Define properly configured WLAN controllers to WCS as follows:

Step 1 Click **Configure > Controllers** to display the All Controllers page.

Step 2 From the command drop-down menu in the right-hand upper corner of the screen, choose **Add Controller** and click **GO,** as shown in Figure 8-3.

Figure 8-3 Adding a Controller to WCS

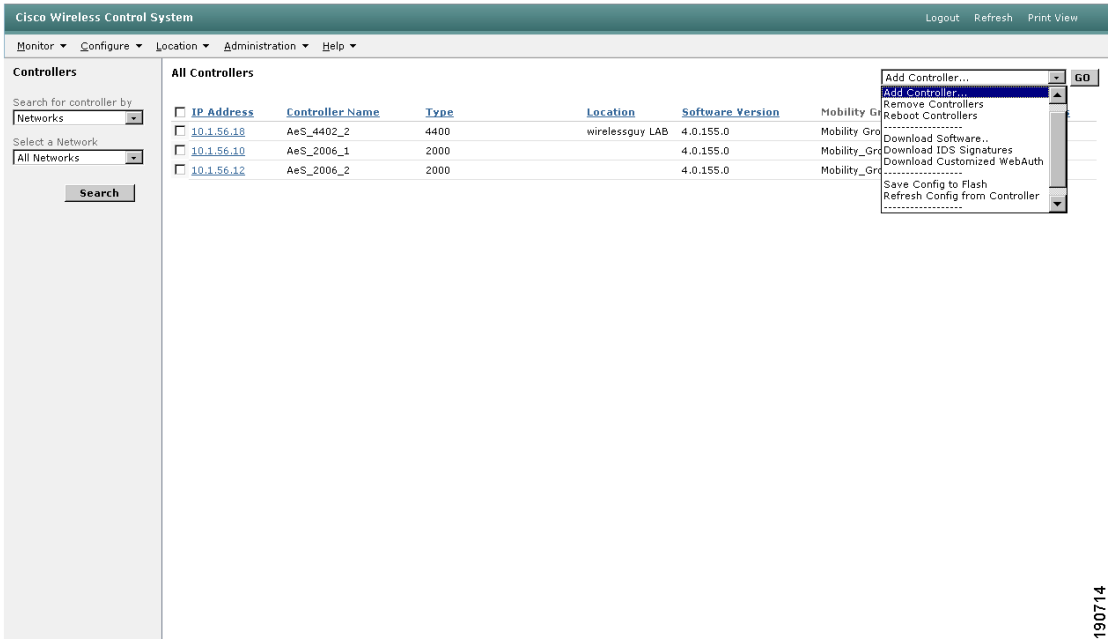

Step 3 On the Add Controller page, enter the controller IP address, network mask and required SNMP settings as shown in Figure 8-4.

Figure 8-4 Defining New Controller IP and SNMP Parameters

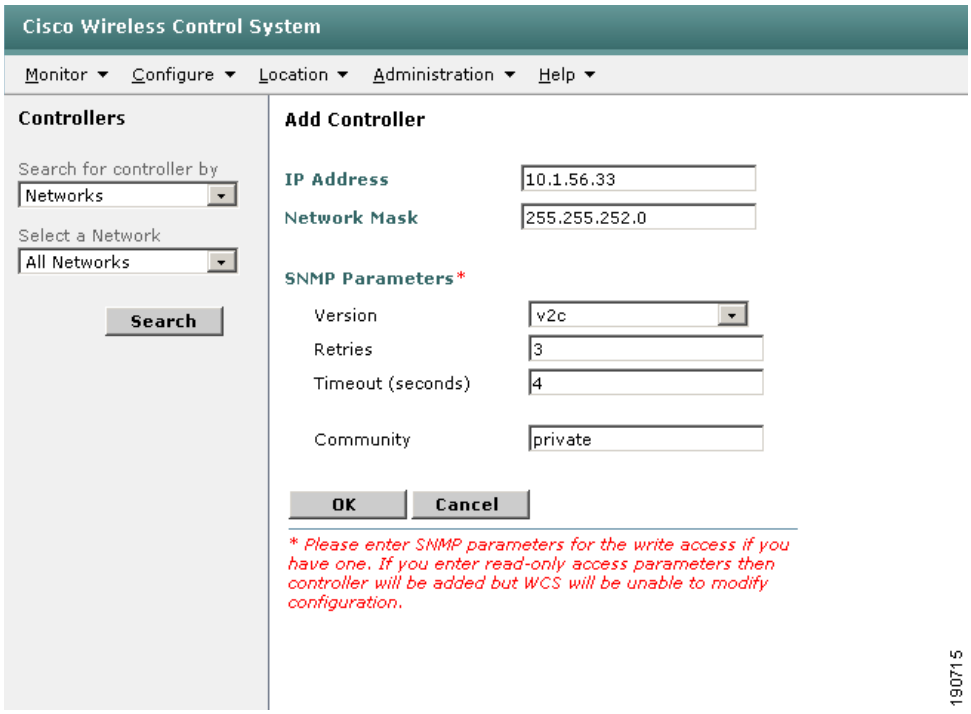

Note that if you use the SNMP read-only community string on this screen, you are able to query (but not modify) controller configurations using WCS. With SNMPv3, read-only access is achieved by specifying the name of a user profile that has been defined in the WLAN controller with an authentication password and privacy password to which read-only access has been permitted. In either case, the use of read-only credentials when attempting to modify a configuration results in a "MIB Access Failed" error message. This occurs whenever WCS attempts to modify the value of a parameter in a controller but SNMP read-only access has been specified. Note that WCS still modifies its internal database with the change even though the WLAN controller itself could not be modified because of the read-only community string. Therefore, if you receive a "MIB Access Failed" message, be sure to back out any changes made to the WCS database, either manually or by using the selective or non-selective synchronization methods described in Synchronizing WCS with Controller and Access Point Configurations, page 8-34.

Step 4 Click **OK**.

WCS displays a "Please Wait" dialog box while it contacts the controller and adds the current controller configuration to the WCS database. It then returns you to the Add Controller page.

- If WCS does not find a controller at the IP address that you entered for the controller, the Discovery Status dialog displays this message: "No response from device, check SNMP communities, version or network for issues". Check these settings to correct the problem:
- The controller port IP address might be incorrect. Check the port setting on the controller.
- WCS might not have been able to contact the controller. Make sure that you can ping the controller from the WCS server operating system.
- The SNMP settings on the controller might not match the SNMP settings that you entered in WCS. To verify this, login to the controller using the web interface or the CLI and make the appropriate corrections as directed in the *Cisco Wireless Control System Configuration Guide, Release 4.0*.

Step 5 Add additional WLAN controllers by repeating these steps if desired.

Restricting SNMP v1/v2c Access using Source IP Address

SNMPv1 and SNMPv2c access to WLAN controllers is typically limited on the basis of whether the management system has been configured with the correct read-only or read-write community strings for the device. Cisco WLAN controllers allow you to add another layer of SNMP access restriction using source IP addresses as well. This makes it more difficult for an unauthorized SNMP manager that somehow has obtained your community strings from gaining control of your WLAN controllers (keep in mind that SNMP v1 and v2c community strings are sent in the clear and can easily be seen using an Ethernet protocol analyzer, as seen in Figure 8-5, where the community string is use is “private”.

Figure 8-5 SNMPv2c Ethereal Trace Showing Plainly Visible Community Strings

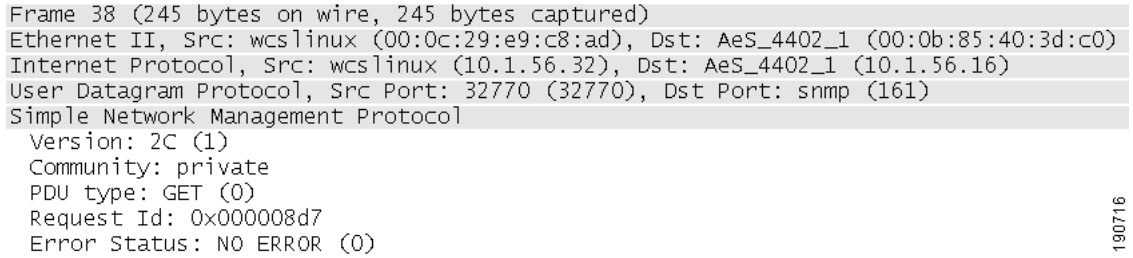

By associating valid IP addresses (or a range of addresses) with each defined community string, only SNMP v1/v2c commands coming from these sources addresses are honored by WLAN controllers so configured, even if the correct community strings are specified.

To configure your controllers in this fashion, use the **Management > Communities** menu option in the controller web interface (*not* WCS) and add IP address and netmask information to your community string definitions. You can also perform this via the controller CLI by using the **config snmp community ipaddr** ***ip-address ip-mask name*** command.

This source address-based restriction capability is not used with WLAN controllers using SNMPv3. SNMPv3 does *not* use community strings and sends all SNMP Protocol Data Units (PDUs) encrypted between WCS and the WLAN controllers. Figure 8-6 shows an example of a protocol analyzer trace of an encrypted SNMPv3 PDU.

Figure 8-6 Example of Encrypted SNMPv3 PDU

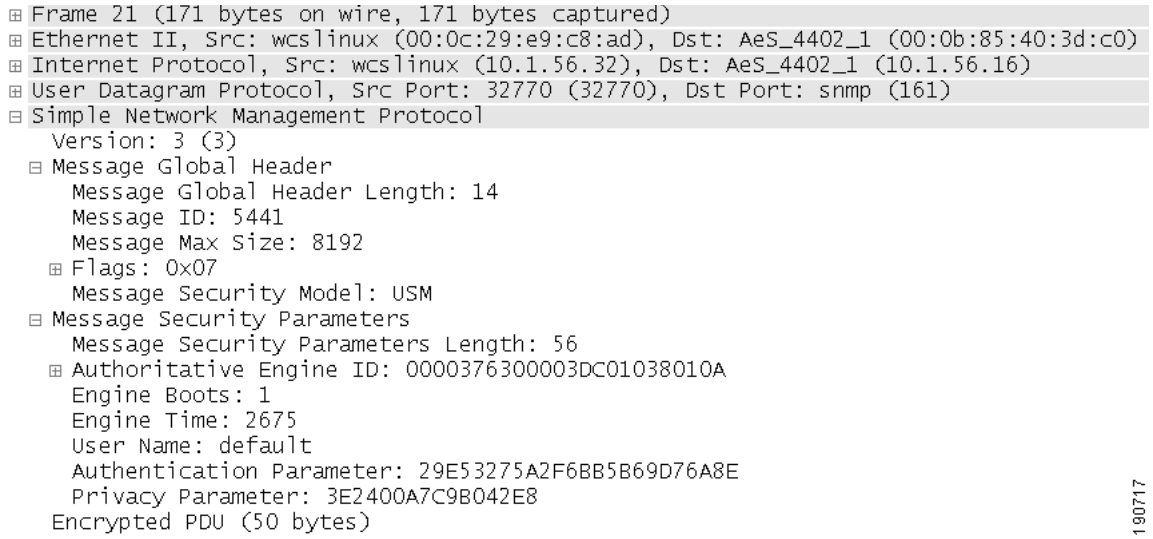

Further information about this capability can be found in the *Cisco Wireless Control System Configuration Guide, Release 4.0*.

Adding Location Appliances To WCS

The Cisco Wireless Location Appliance enhances the capabilities of a location-enabled WCS server by computing, collecting, and storing historical location data and allowing WCS to display graphical location information for multiple clients, tags, and rogue devices simultaneously.

Configuration of the location appliance is performed from WCS using the menus and submenus located under the main menu **Location** tab after initial configuration of IP parameter settings, as described in the *Cisco Wireless Location Appliance—Installation Guide,* available at the following URL: http://www.cisco.com/en/US/products/ps6386/products_installation_and_configuration_guide_book09186a00804fa761.html.

To define a location server(s) to WCS, follow these steps:

Step 1 Click **Location > Location Servers** to display the All Location Servers page.

Step 2 From the command drop-down menu in the right-hand upper corner of the screen, choose **Add Server** and click **GO.**

Step 3 Enter the required information as shown in Figure 8-7.

Figure 8-7 Defining a Location Appliance to WCS

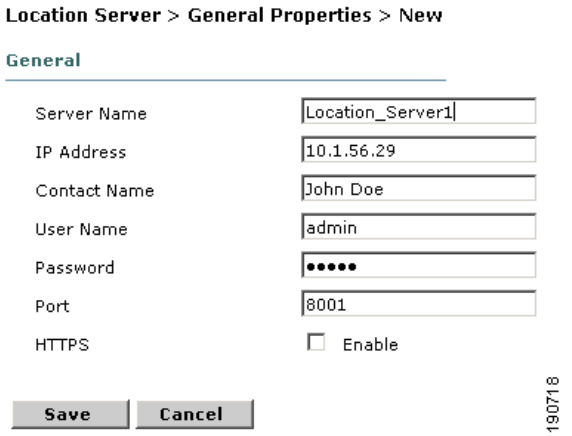

Step 4 If you want to enable HTTPS, enable the check box only *after* completing steps 1 through 3 completely. After enabling the check box, click **GO** again.

Using WCS to Configure Your Wireless Network

Configuring Network Components

After network components have been successfully defined to WCS and two-way communication via SNMP has been established, these devices can be configured and managed centrally as part of the management domain of that WCS server. Of course, WCS allows devices to be configured one parameter screen at a time in a similar fashion to the controller web interface that is available when accessing the WLAN controller individually. However, WCS goes much further and allows for the provisioning of *policy templates* that can be applied to WLAN controllers and lightweight access points. Policy templates are groups of configuration parameters that in most cases are defined once and then applied to multiple controllers without the need to manually re-key each value and send each screen of configuration data to each controller and lightweight access point individually. After being defined and implemented in WCS, the use of policy templates greatly reduces the possibility of controller misconfiguration by ensuring that the proper values are defined once and then saved for future re-application.

When using the commands in the following subsections allowing multiple target objects to be selected, these objects *must* all be present on one display page. This is important if the total population of controllers displayed, for example, spans several pages and requires paging forward and backward.

Configuring WLAN Controllers

WCS allows for the configuration of network components via the **Configure** option on the main menu bar. WLAN controllers can be configured via **Configure > Controllers,** and lightweight access points can be configured via **Configure > Access Points**. In this manner, controller and lightweight access points can be either individually configured or the configuration that was applied to a group of devices via the application of a policy template can be overridden.

To configure a WLAN controller that is currently managed by WCS, perform the following steps:

Step 1 Click **Configure > Controllers** to display the All Controllers page. In large networks, you may find it helpful to use the "Search for Controllers" filter in the left-hand column to narrow the selection of displayed controllers by name, IP address, or network. In large networks, the listing of controllers can be sorted in ascending or descending order by clicking on the appropriate column heading.

Step 2 Click on the hyperlink representing the IP address of the WLAN controller that you want to configure. Note that configuration of a device cannot be performed by simply enabling the check box for the controller; the hyperlink for the particular WLAN controller must be used.

Step 3 On the Controller Properties screen, you may change the name assigned to this WLAN controller as well as the location text string and the controller SNMP properties.

There are also three check box fields available on the **Controller Properties** screen:

- Restore on Cold Start Trap—When this check box is enabled, WCS initiates the **Restore Config to Controller** function on reception of a SNMP cold-start trap, indicating that a WLAN controller has rebooted.

Note For further information, see section 13.7 CSCsc59232 —4400 and 2006 Controllers Not Issuing Cold Start Traps.

This procedure entails WCS refreshing the configuration in the controller from the current contents of the WCS database. This is a valuable feature designed to ensure that the configuration loaded into a freshly-booted WLAN controller is indeed the configuration of record currently contained within the WCS database. In this manner, any unauthorized local changes made to the controller configuration via the controller web interface or CLI are overridden with the configuration of record stored in the WCS database.

Note that the configuration programmed into the controller is not implicitly saved when the restore on cold-start feature is used. This means that the controller retains its original configuration in nonvolatile memory and not the changes that were transferred to it in conjunction with the cold-start restore. If this is not desired, after the controller is fully booted and has received its configuration from WCS, perform an explicit save of the running configuration of the controller to nonvolatile (flash) memory using the Save Config to Flash function as shown in Figure 8-10.

See Non-Selective Synchronization, page 8-36, for further details on the **Restore Config to Controller** function.

- Refresh on Save Config Trap—When this check box is enabled, WCS initiates the **Refresh Config from Controller** function upon reception of a save-config trap (bsnConfigSaved) indicating that the current configuration in the WLAN controller has been saved to the controller nonvolatile (flash) memory. WCS then refreshes the configuration contained in its databases with the current configuration of the controller. Any configuration objects found in WCS but not found in the controller configuration are retained in the WCS databases. See Non-Selective Synchronization, page 8-36 for further details on the **Refresh Config from Controller** function.
- Save Before Backup—This check box has an effect only when the **Configuration Backup** scheduled task has been enabled and submitted for execution. (An identical but independent check box appears for **Configure > Controllers > *controllerIPadddress* > System > Commands > Upload/Download Commands > Upload/Download Commands > Upload File from Controller**.)

  When enabled, **Save Before Backup** indicates that the running configuration of this controller should be saved to the nonvolatile memory of the controller before the scheduled task archiving the controller-saved configuration. Because the **Configuration Backup** scheduled task archives only the saved configuration of the controller and not the currently running configuration, enabling this

check box to ensure that any recent unsaved changes are saved and therefore included in the archive. See Configuration Backup, page 8-117 for further details on the **Configuration Backup** scheduled task.

Step 4 You may now select from the list of configuration object categories listed in the column on the left-hand side of the **Controller Properties** screen as shown in Figure 8-8. Guidance on configuring the parameters contained in each of the controller configuration categories can be found in the WCS main menu bar under **Help > Online Help**.

Figure 8-8 Controller Properties

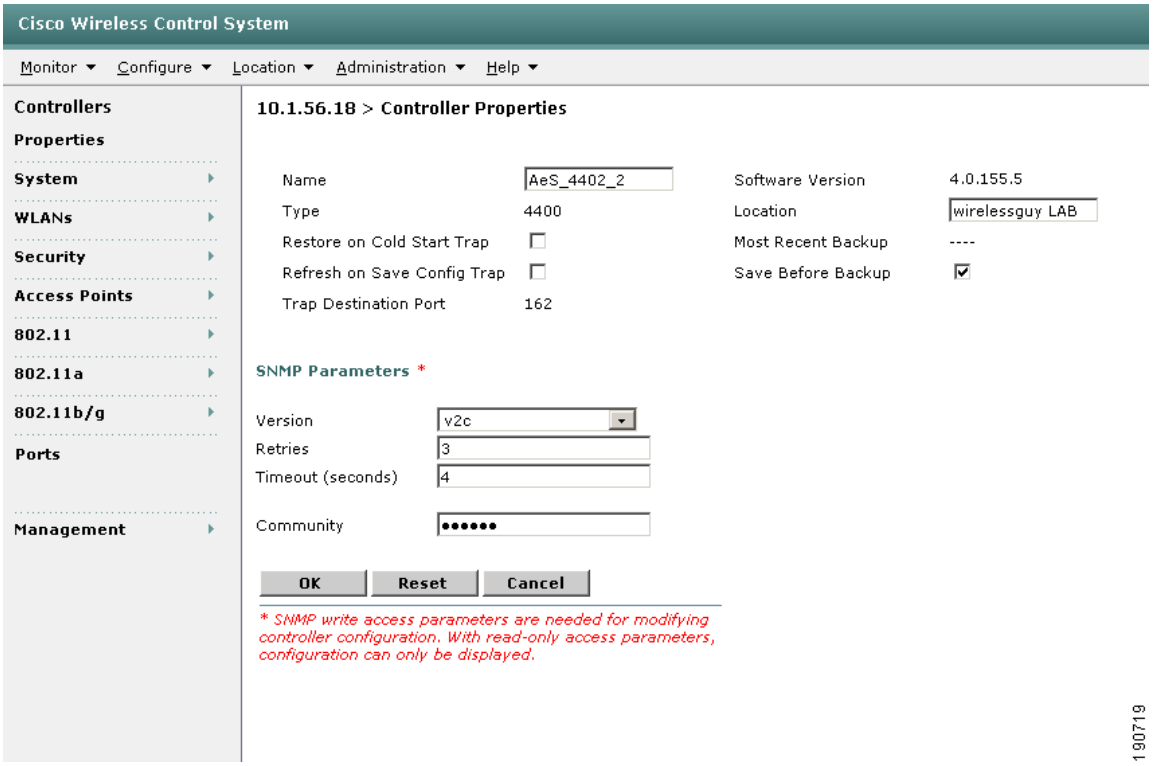

After making your desired changes in each of the various controller configuration object categories, you need to save your changes (as shown in Figure 8-9) **in each category** for your changes to be applied to the current running configuration of the controller.

Figure 8-9 Save and Apply Changes for Configuration Object Category "Network Time Protocol"

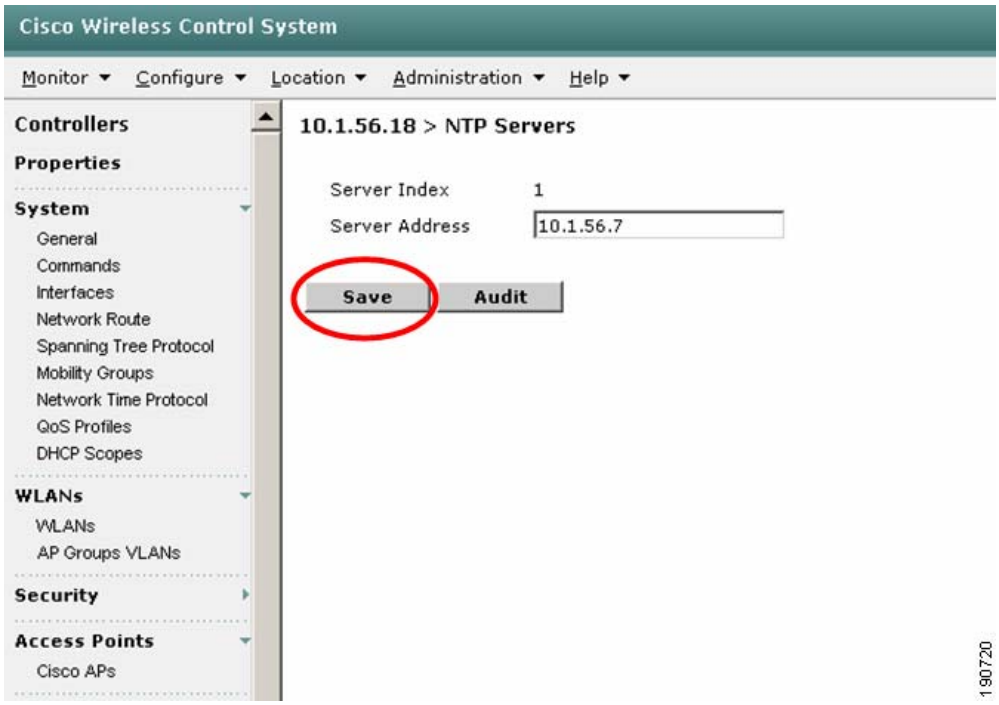

Keep in mind that the procedure outlined thus far *does not* save your changes to the nonvolatile memory of the controller (that is, your changes are lost if the controller is rebooted or loses power). To write your newly modified controller running configuration to nonvolatile memory, perform the following steps:

Step 1 Click **Configure > Controllers** to display the All Controllers page.

Step 2 Select the check box(es) for the controller(s) for which you want to write the running configuration to nonvolatile memory.

Step 3 From the command drop-down menu in the right-hand upper corner of the screen, choose **Save Config to Flash** (shown in Figure 8-10) and click **GO.**

The running configuration for each controller selected is written to their respective nonvolatile memories.

Figure 8-10 Saving Configuration to Controller Nonvolatile (Flash) Memory

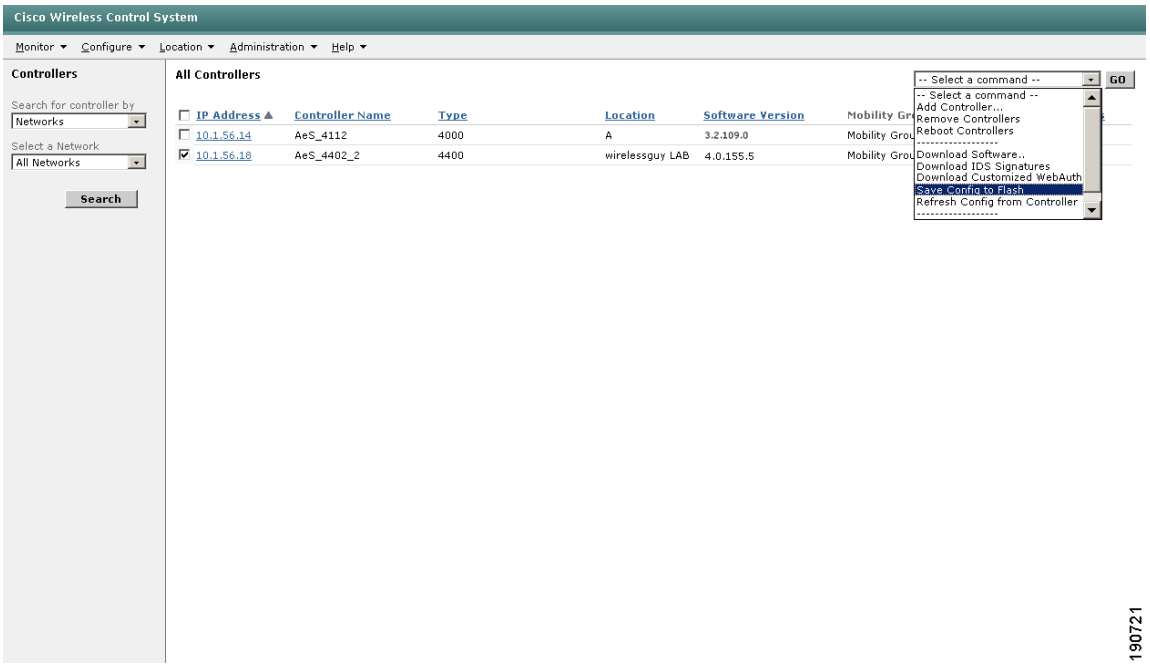

Further guidance about configuring WLAN controllers can be found in the WCS main menu bar under **Help > Online Help**.

Configuring Lightweight Access Points

Lightweight access points can be configured using WCS in a similar fashion to that described in the previous section on WLAN controllers. As mentioned previously, WCS does not configure lightweight access points directly but rather does so via the SNMP agent and other software components present in the WLAN controller to which the lightweight access points are currently registered. Thus, only access points that are registered with controllers can ultimately be managed via WCS.

By using the **Configure > Access Points** menu option, lightweight access points can be individually configured or the configuration that was applied to a group of access points using policy templates can be overridden.

Step 1 Click **Configure > Access Points** to display the All Access Points page.

In large networks, it is helpful to narrow the listing by using the "Search for APs By" filter in the left-hand margin of the screen (see Figure 8-11). Access points can be filtered based on several filter types such as MAC addresses, AP name, assigned controller, unassociated or unassigned status and outdoor, campus, building, or floor location.

Figure 8-11 All Access Points Display Menu

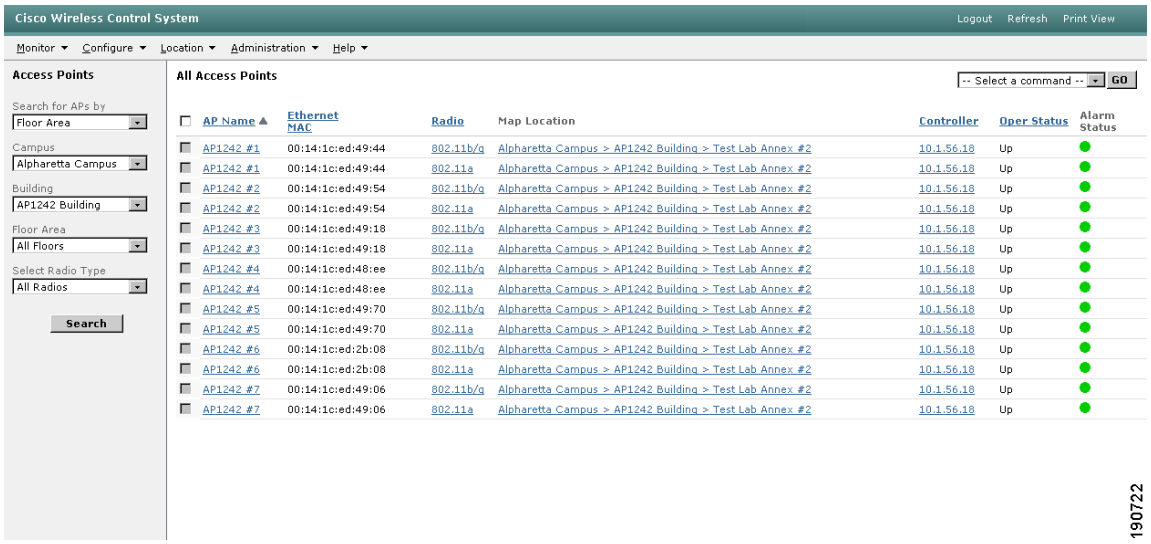

Note that the All Access Points page (shown in Figure 8-11) is a bit more involved than the All Controllers page seen in previous sections. Each dual-band lightweight access point is actually represented twice in WCS with a separate line item entry for each radio interface contained within the lightweight access point. Therefore, unless a specific radio type is selected using the display filter in the left-hand column, a typical dual-band lightweight access point has two line item entries on the "All Access Points" menu. Each entry is differentiated by the values listed under the Radio column heading.

Note also that unlike in All Controllers, there is more than just a single hyperlink entry per line item in the All Access Points menu. Clicking on the AP name takes you to the general lightweight access point configuration panel, while the Radio identifier takes you directly to the submenu for a specific radio in that lightweight access point. Clicking on the Location hyperlink immediately links you to the location map where that lightweight access point has been assigned. Clicking on the Controller hyperlink takes you to the Controller Summary screen for the WLAN controller to which this lightweight access point is assigned.

Figure 8-11 shows the All Access Points menu sorted by campus, building, and floor. Access points that are currently registered with controllers show the controller IP address as a hyperlink in the Controller column. Attempting to configure an access point that is not registered results in the error "This AP is not associated with any Controller" being displayed.

You are now ready to select an access point that is registered with a controller and to modify its configuration.

Step 2 Select the desired registered lightweight access point from the All Access Points menu by clicking on the AP name hyperlink.

Step 3 The **Access Point >** ***ap name*** screen is now displayed, as shown in Figure 8-12.

Figure 8-12 Access Point > ap name Configuration Screen

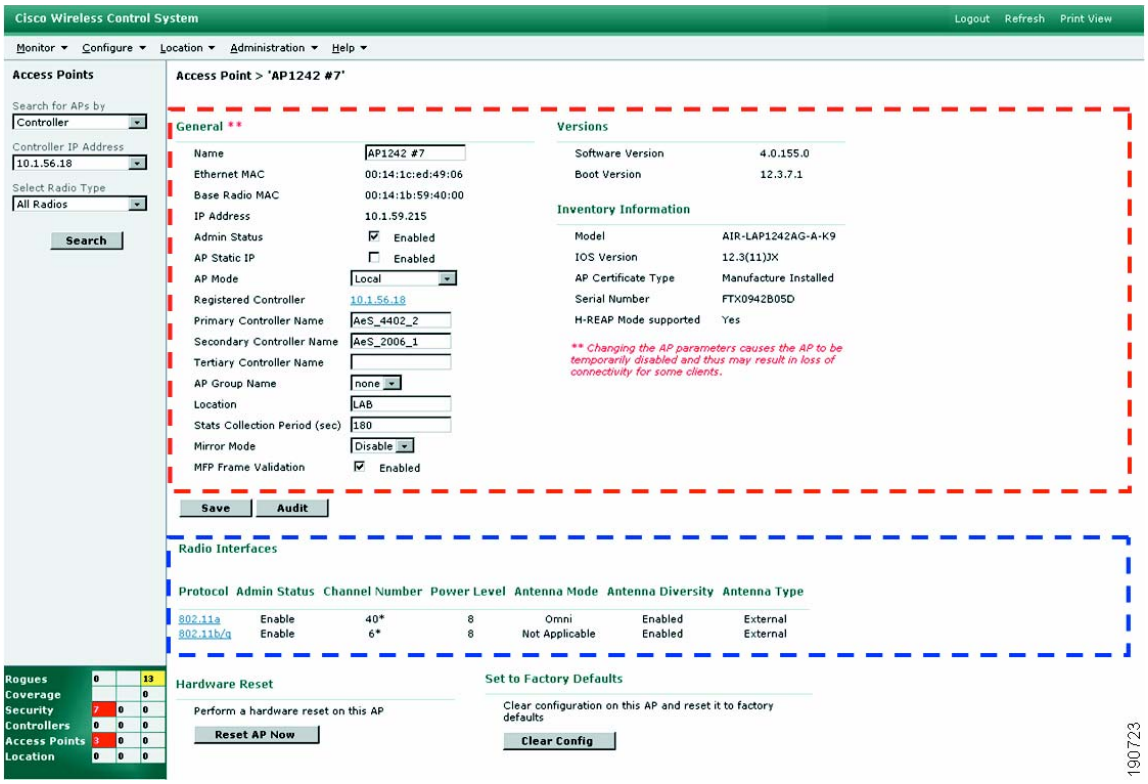

Any parameters modified within the confines of the red dashed rectangle indicate areas that apply to the access point in general. The “Name” parameter shown in this red dashed rectangle is default to the value “AP” concatenated with the MAC address of the access point. You may find it useful to use this field to assign a new name that conveys more meaning within the context of your particular organization, or perhaps to assign a numerical differentiator to each name that assists when sorting and searching the full list of managed access points (seen in Figure 8-12). An example of how this can be used can be seen in Figure 8-11, where access points are named in numerical sequence (“AP1242#1”, “AP1242#2”, and so on) sorted by name in numerical order.

The two radio protocol hyperlinks under the radio interfaces heading (located within the blue dashed rectangle) provide access to **Access Point > *ap name* > 802.11a** and **Access Point > *ap name* > 802.11b/g** radio specific sections (shown in Figure 8-13). Changes to parameters contained within these radio-specific screens affect only the radio interface concerned.

Figure 8-13 Radio-Specific Access Point Configuration

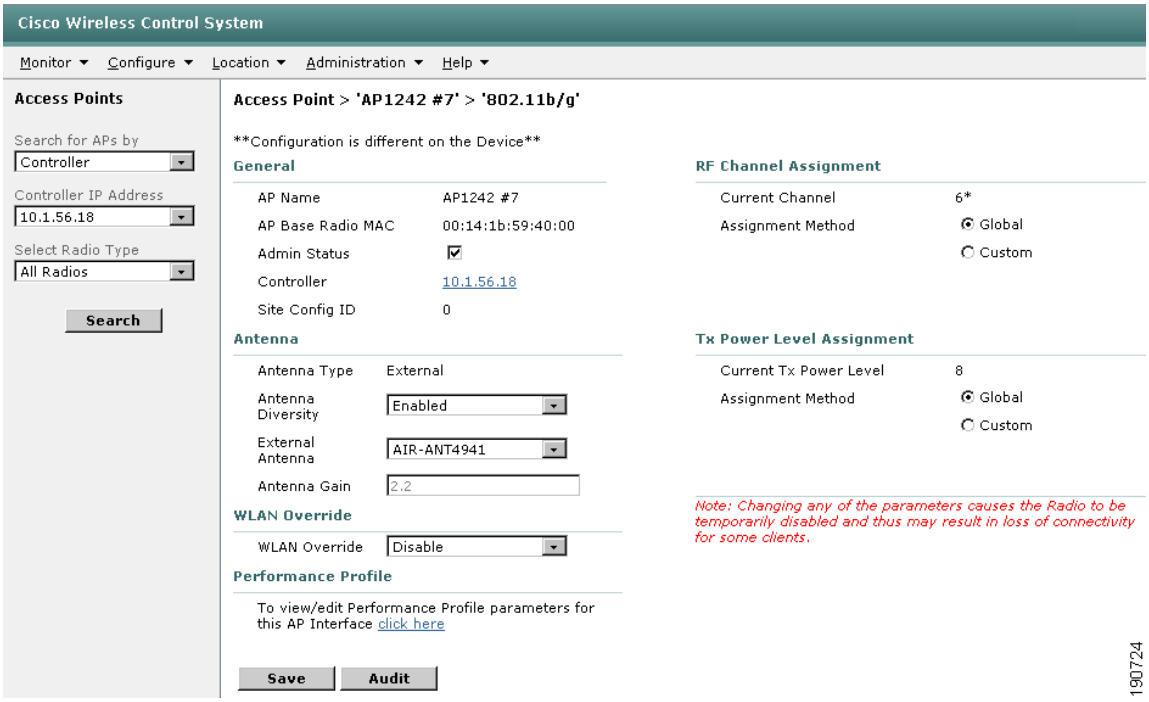

Note that Figure 8-12 also provides a few other options. For example, you find the ability to issue a hardware reset only on the access point or performing a hardware reset and setting the access point configuration to factory defaults. In addition, an option is present to audit configuration parameters. This audit option compares the settings stored in the WCS database to those that are currently resident in the lightweight access point. If they differ, you are presented with a screen similar to that shown in Figure 8-14 asking which set of values (those contained within WCS databases or those contained within the access point/controller) should prevail.

Keep in mind that any parameters changed on this page must be saved to the WCS database (using the **Save** button) before being subject to comparison as part of an audit. If entries are changed and the Audit button is used *before* the changes are saved, the changed entries are discarded. In addition, when an Audit button appears in a configuration-section specific menu such as this, *only* the values contained in the WCS database for the parameters shown on the page are subject to the audit. You are not notified of any discrepancies between the current AP/controller configuration and stored WCS values for parameters other than those shown here.

Figure 8-14 Access Point Audit Report

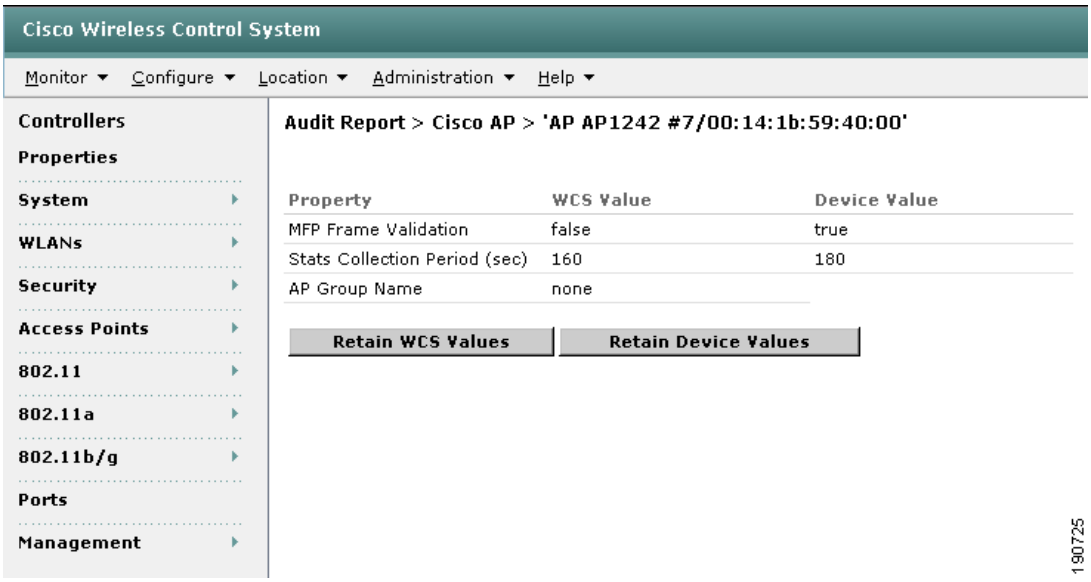

Step 4 When you are satisfied with your changes on each page, click **Save** and the values are written to both the WCS database as well as the AP/controller configuration.

Configuring WLAN Controllers, page 8-12, showed that when configuring WLAN controllers with WCS, the new configuration is applied as a running configuration and is not automatically saved to nonvolatile memory. However, this is not the case when configuring lightweight access points. When configuring lightweight access points, any changes that are applied from WCS to the access points via the controller are saved to the nonvolatile memory of the access points. Therefore, there is no need for an explicit save procedure to ensure that your changes are still intact after access points are rebooted. In fact, after your changes are applied, they migrate with lightweight access point even if the lightweight access point become registered to a different controller.

Further guidance about configuring lightweight access points can be found in the WCS main menu bar under **Help > Online Help**.

Copying Lightweight Access Point Configurations

In some cases, it may be necessary to copy the configuration that is stored in the WCS database for a lightweight access point to a new lightweight access point. A good example of when this might be necessary is when replacing a lightweight access point that has become damaged in some way with a replacement lightweight access point that has been sent via the Cisco SmartNet program.

WCS makes it possible to copy the configuration-of-record stored in WCS for the original (source) access point and apply it to the replacement (target) access point. When performed, the configuration of the target lightweight access point is overwritten and it assumes the majority of the parameters originally configured for source lightweight access point.

Note Admin Status, Monitor Mode, WLAN Overrides, Channel Assignment, and Antenna Diversity settings are not copied.

Only lightweight access points that are known to WCS but not registered with a WLAN controller can serve as the source of the lightweight access point copy operation. Similarly, only lightweight access points that are known to WCS and currently registered with a controller within this management domain can serve as the target of a lightweight access point copy operation. From the perspective of WCS, it does not matter if the source lightweight access point was originally used on a different WLAN controller than where the target lightweight access point is installed.

To perform this copy operation, follow these steps:

Step 1 Click **Configure > Access Points** to display the All Access Points page. In large networks, it is helpful to narrow the listing by using the "Search for APs By" filter in the left-hand column margin.

Step 2 Select the non-registered source lightweight access point that you want to copy the configuration from by enabling the check box next to its name (notice that the checkboxes for both radio interface line items become enabled).

Note A *non-registered* access point is an access point that had previously been registered to a WLAN controller defined to WCS but is not currently registered to any WLAN controller in this management domain.

Step 3 From the command drop-down menu in the right-hand upper corner of the screen, choose **Copy and Replace AP** and click **GO.**

Step 4 Select the registered target lightweight access point that you want to have serve as the destination for the copied configuration. Enable the check box for "Copy Location Information" if you want the location map information for the source lightweight access point copied as well (that is, this positions the target lightweight access point to the same coordinates on location floor maps as the source lightweight access point).

Step 5 Click **Copy To AP** to copy the configuration. The current configuration of the target lightweight access point is replaced with the configuration of the source lightweight access point.

The copy and replace operation is now complete. Notice that the destination registered AP is now configured with the name of the originating non-registered lightweight access point. To avoid confusion, the name of the source or target lightweight access point should be changed or the source lightweight access point deleted, as described in Removing Lightweight Access Point Configurations, page 8-21.

Additional information on copying lightweight access point configurations can be found in the WCS main menu bar under **Help > Online Help**.

Removing Lightweight Access Point Configurations

After the **Copy and Replace AP** operation in Copying Lightweight Access Point Configurations, page 8-20 is performed, you are still left with the original lightweight access point definition resident in WCS. In the case of the replacement of a damaged lightweight access point, the source definition is no longer used because the replacement lightweight access point has assumed its duties. In this case, Cisco recommends that after the lightweight access point configuration has been copied, the original access point configuration should be removed. Otherwise, you will have the old unused access point configurations simply cluttering up the WCS database and adding unnecessary confusion to WCS screens when displaying lists of all lightweight access points.

You can use the Remove APs command in WCS to remove the configuration for the original lightweight access point. Keep in mind that *only* lightweight access points that are not currently registered with any WLAN controller can be removed from the WCS database via the Remove AP operation.

To remove a non-registered lightweight access point configuration from WCS, perform the following steps:

Step 1 Click **Configure > Access Points** to display the All Access Points page. In large networks, it is helpful to narrow the listing by using the "Search for APs By" filter in the left-hand margin.

Step 2 Select the non-registered access point(s) that you want to remove by enabling the appropriate check box(es).

Step 3 From the command drop-down menu in the right-hand upper corner of the screen, choose **Remove APs** and click **GO.**

Step 4 Confirm your intention to remove the lightweight access points. After doing so, the selected lightweight access point(s) are removed from WCS.

Additional information on removing lightweight access point configurations can be found in the WCS main menu bar under **Help > Online Help**.

Defining and Applying Policy Templates

Policy templates are groups of configuration objects that are typically defined once and then applied to multiple controllers without the need to manually re-key each object value and send each screen of configuration data to each controller, lightweight access point, or radio interface individually. After being defined and implemented in WCS, the use of policy templates greatly reduces the possibility of controller misconfiguration by ensuring that the proper values are defined once and saved for future use when defining subsequent resources.

Policy templates allow for the creation of configuration objects along with a simple means with which to propagate those configuration objects among multiple WLAN controllers, lightweight access points, or access point radios. By using WCS policy templates, uniform QoS, security, and RF management policies can be easily created and enforced across an entire enterprise or outdoor deployment.

The definition of WLAN controller and lightweight access point policy templates is a relatively straightforward task and is similar in many aspects to the procedure described in Configuring WLAN Controllers, page 8-12 and Configuring Lightweight Access Points, page 8-16 for directly configuring managed resources.

Complete guidance concerning how to properly define policy templates within WCS can be found in the *Cisco Wireless LAN Controller Configuration Guide, Release 4.0* and also in the WCS main menu bar under **Help > Online Help**. In addition, the following are key points to keep in mind when planning to use WCS policy templates to assist in managing your wireless LAN enterprise network:

- Policy templates can be applied to WLAN controllers, lightweight access points, and their radio interfaces but not to location appliances.
- Policy templates can be explicitly applied ("pushed") to WLAN controllers, access points, and radios.
- Changes applied to network resources via policy templates can be overridden by authorized operators via WCS or the local controller GUI/CLI interface. The use of the Restore on Cold Start Trap option described in Configuring WLAN Controllers, page 8-12 ensures that the WLAN controller is restored to the WCS configuration of record whenever the controller is rebooted.

Note For further information, see section 13.7 CSCsc59232—4400 and 2006 Controllers Not Issuing Cold Start Traps.

This stored configuration is updated whenever a policy template is successfully applied to a managed resource by WCS.

- Policy templates can be applied to more than a single device at once, making configuration of multiple controllers or lightweight access points easy and efficient. The **Configure > Config Groups** option enables the grouping of multiple templates for application to one or more controllers within the same mobility group (see Using Policy Template Configuration Groups, page 8-25).
- When a controller policy template is created and it is:
  - Saved—This saves the policy template in the WCS database as an unapplied policy template. The template can be applied to managed resources now or at a later time. If you make changes to a policy template but do not save the template or attempt to apply it to at least one controller, the changes are not available on the next use of the policy template. (The policy template is implicitly saved as soon as the **Apply to Controllers** button is clicked, even if no controllers are then subsequently selected for application.)
  - Applied to controllers—This issues an implicit save of the template and applies the policy template to at least one controller. If the application of the template to the controller(s) is successful, the stored configuration for that controller(s) is updated in the WCS database as well.
- Policy templates allow for the definition of the most commonly defined configuration objects (see Figure 8-15 for a listing of available controller policy template configuration object categories).

Figure 8-15 Configuration Object Categories Available Via Policy Templates

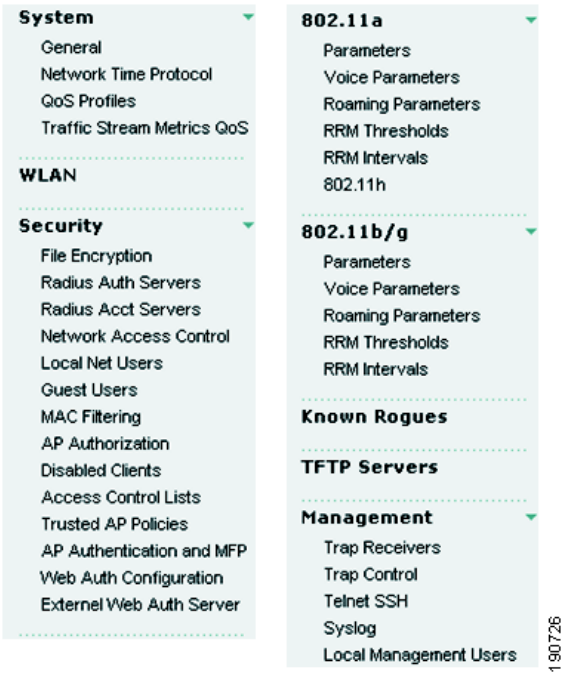

For some configuration parameters, explicit configuration is required via the WCS **Configure > Controllers** facility. (Some WLAN controller parameters must be configured via the controller web interface or the CLI. Examples of this include SNMP trap destination port, NTP polling interval, and serial port configuration.) The majority of the configuration objects not addressed by policy templates are typically site or controller unique, which tends to exclude them from application as part of a enterprise-wide policy template.

- If you wish to save changes enacted by the application of policy templates in the non-volatile saved configuration of a controller, the save configuration function should be explicitly performed for the controller or group of controllers after the policy templates have been applied. Policy templates are not automatically re-applied to network components after they are re-booted and become reachable from WCS.
- Access points must be registered to WLAN controllers to be eligible to have access point/radio policy templates applied to them via **Configure > Access Point Templates** (shown in Figure 8-16).

Figure 8-16 Access Point/Radio Policy Template

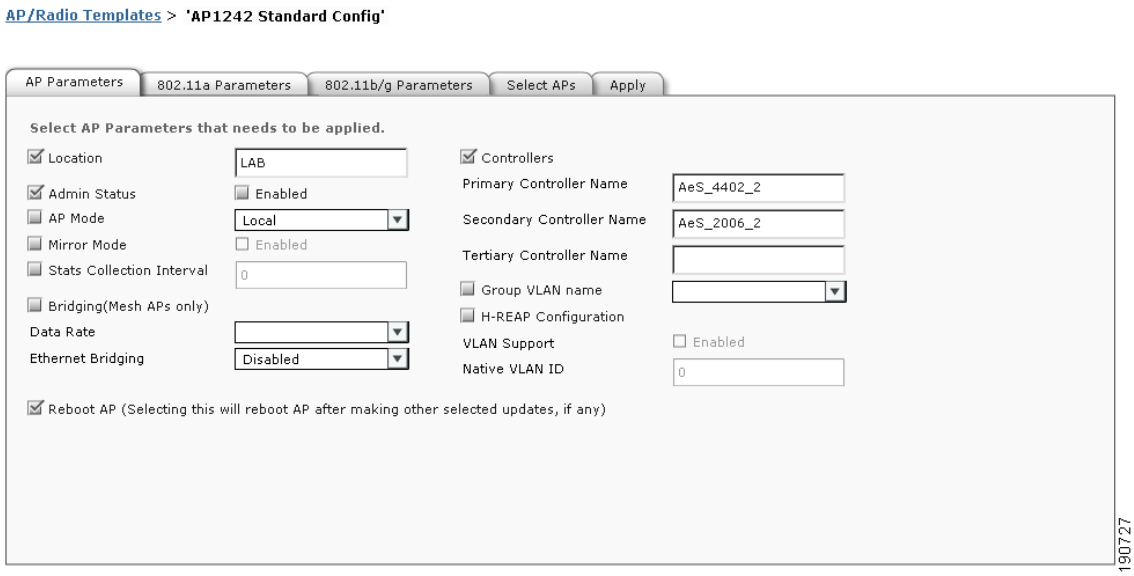

These access points can be on a single controller or spread among two or more controllers (the search parameters in the left-hand margin of Figure 8-17 facilitate choosing access points). Note that after selecting the parameters you wish to configure in all configuration tab areas, you must save the template before applying it to any access points (see red circle in Figure 8-17). Failure to save the template before application results in any changes made being ignored. Beginning with release 4.0 of WCS, access point and radio templates can be saved in WCS for subsequent re-use.

Figure 8-17 Saving and Applying the AP/Radio Template

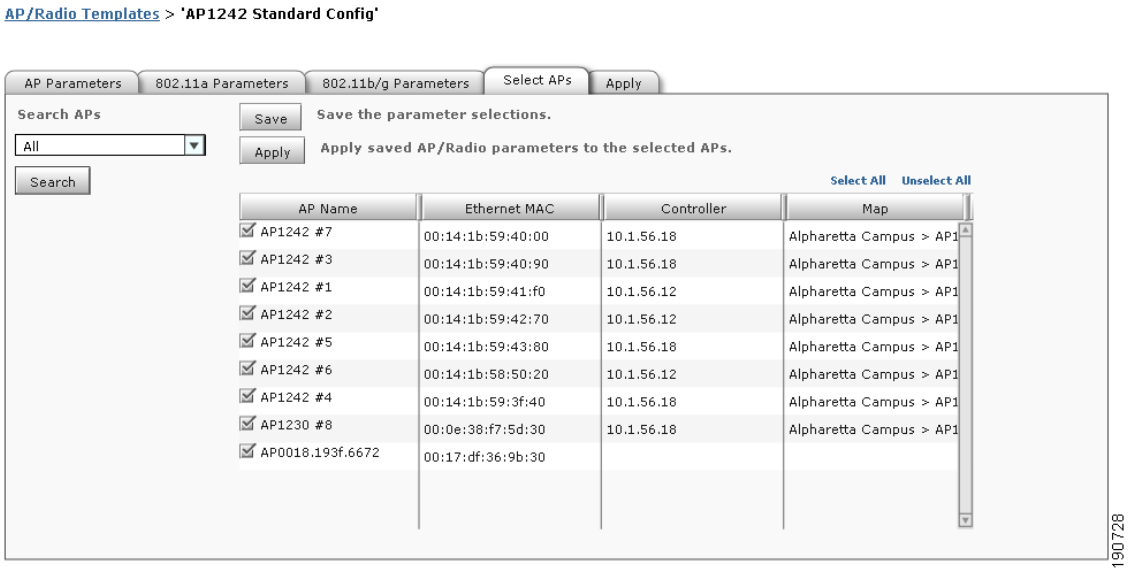

- After being applied, lightweight access point and radio policy templates are automatically saved in lightweight access points. Lightweight access point and radio configuration changes that have been applied via policy templates are available on a controller or access point re-boot, because the values are saved in the nonvolatile (flash) memory of the lightweight access point. If a controller should fail and the lightweight access point migrate and register to an adjacent controller, the changes that have been applied via the policy template remain with that lightweight access point or its radio interfaces unless changed by policies on the new controller.
- When defining access point and radio policy templates, only configuration objects whose checkboxes are enabled are transmitted to the device. Any configuration objects already existing in the lightweight access point or radio interface are retained if their associated value in the template is not specified. If you want to remove or "blank out" an existing parameter, the configuration object in the access point or radio template must have its check box enabled and the appropriate blank value specified.

Using Policy Template Configuration Groups

By creating a configuration (config) group, you can group controllers that should have the same mobility group name and similar configuration. You can assign templates to the group and push templates to all the controllers in a group. You can add, delete, or remove config groups, and download software, IDS signatures, or a customized web authentication page to controllers in the selected config groups. You can also save the current configuration to the nonvolatile (flash) memory of all controllers in selected config groups.

Complete guidance on the use of Configuration > Config Groups can be found in the *Cisco Wireless Control System Configuration Guide, Release 4.0*.

Keep the following points in mind when using policy template configuration groups:

- The mobility group already assigned to a controller is changed by the function of the configuration groups when changes are applied.

- Templates that are applied to controllers via the Config Groups mechanism are not saved to the nonvolatile memory of the controllers by default unless the controllers are rebooted using the Reboot tab, as shown in Figure 8-18.

Figure 8-18 Config Groups Reboot Menu

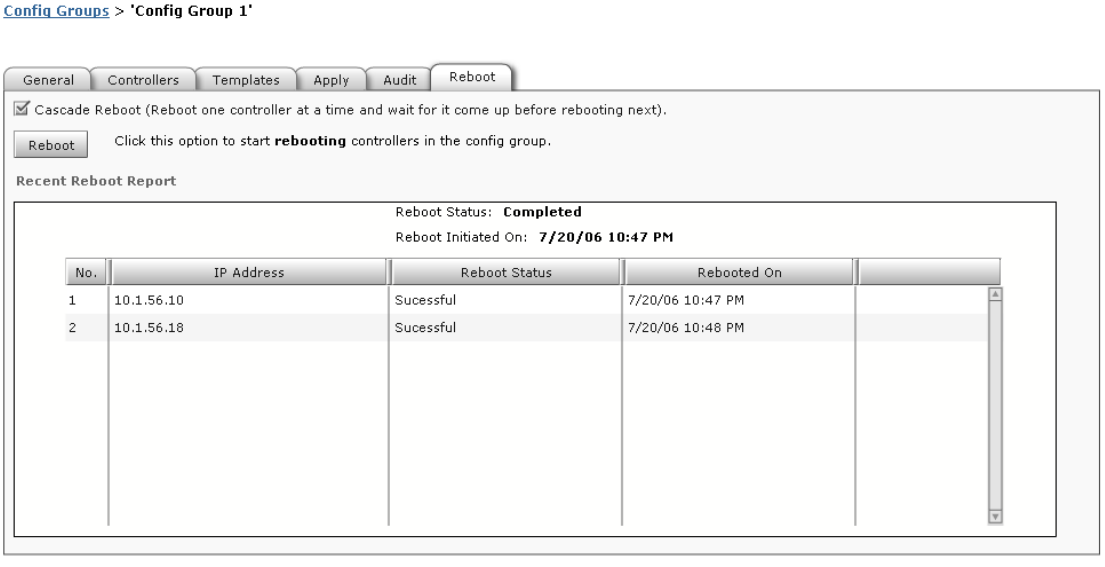

If you wish to save the updated configuration of all controllers in the configuration group to their respective nonvolatile memories without rebooting them, perform the following:

Step 1 From Configure > Config Groups, click the check box(es) to choose one or more config groups on the Config Groups window.

Step 2 Choose **Save Config to Flash** from the Select a command drop-down menu and click **GO.**

You can perform other utility functions on the controllers in the configuration group in a similar fashion, such as downloading controller software, IDS signatures, and web authentication credentials. Complete details about how to perform these tasks and more can be found in the *Cisco Wireless Control System Configuration Guide, Release 4.0*.

Configuring Location Appliances

When a Cisco Wireless Location Appliance is introduced and configured for use within a Cisco Unified Wireless Network that contains a location-enabled WCS server, the location appliance assumes responsibility for several important tasks. Key among these are the execution of location positioning algorithms for multiple devices, the ongoing processing of historical location and statistical information, the issuance of location notifications, and the provisioning of a defined SOAP/XML Application Programming Interface (API) for other business applications wishing to make use of the device positioning information available in the location appliance.

WCS acts in concert with the location appliance by serving as the user interface (UI) for the enhanced services provided by the location appliance. Other than during the initial installation and shutdown, direct user interaction with the location appliance via the CLI is typically not required.

Integrating a Cisco Wireless Location Appliance into a Cisco Unified Wireless Network architecture immediately enables key operational advantages, such as the following:

- Scalability—Adding a location appliance greatly increases the scalability of the Cisco LBS solution from on-demand tracking of a single device to a maximum capacity of 2500 devices (WLAN clients, RFID tags, rogue access points, and rogue clients). For deployments requiring location tracking of greater than 2500 tracked devices, additional location appliances can be deployed as part of the Unified Wireless Network and managed under a common WCS.
- Historical and statistics trending—The appliance records and maintains historical location and statistics information, which are available for viewing via WCS.
- Location notifications—Location-based alarms and notifications can be triggered through area boundary definitions, allowed areas, and distances. These alarms and notifications can also provide advanced warning of rogue movement and appearance/disappearance.
- SOAP/XML API—The location appliance interfaces to WCS using a very rich and robust SOAP/XML API interface. These same capabilities allow for integration with other business applications that can use the location information contained within the location appliance in a variety of creative value-added applications. Asset tracking, inventory management, location-based security, and automated workflow management are just a few examples of this.

Complete guidance about how to use the Location menu option under WCS to properly configure the Cisco Wireless Location Appliance is available in the *Cisco Wireless Location Appliance—Configuration Guide*.

Managing Network Component Software

An often-overlooked but nevertheless critical feature of any effective enterprise network management system is the ability to inventory and update the operating software levels of the various components comprising the network. An effective enterprise wireless network management system must be able to regularly inventory software levels and facilitate their upgrade by authorized personnel from either centralized or distributed software repositories.

It is generally regarded as standard industry best practice for network architects and network management staff to be aware of current operating software levels throughout the network. Periodic reviews of http://www.cisco.com and regular discussions with your Cisco account team should be conducted to keep abreast of new features, feature improvements, and bug fixes as they become available and posted. Software updates should be applied to your network only after a careful analysis of new enhancements and bug fixes has been performed and a determination made of the applicability of these software updates to your specific environment. This may be done in conjunction with your local Cisco account systems engineering representative or the Cisco Technical Assistance Center.

WCS offers the ability to effectively manage device operating software such as device operating systems, web certificates, and IDS signatures across various components of the Cisco Unified Wireless Network. In addition, WCS makes it possible to routinely archive controller configurations (as well as the WCS database itself) to protect against potential inadvertent loss of data. The subsections that follow describe how WCS provides effective management of these categories of operating software in wireless LAN controllers, lightweight access points, and location services appliances.

Keep in mind that the level of operating software in registered lightweight access points is automatically managed by the WLAN controller. The version of operating software loaded into any registered lightweight access points is dependent on the level of operating software present in the controller. Although WCS clearly displays the current software levels in each lightweight access point, there is no need (and therefore no ability exists in WCS) to explicitly manage the level of operating software present in lightweight access points.

Managing Controller Operating Software, Web Authentication Bundles, and IDS Signatures

Cisco WCS allows for WLAN controller operating software to be updated via two approaches. Each of these involves the use of the TFTP file transfer protocol, but there are differences in the source server that is used to update the controller.

The actual data transfer can be configured to occur between one of the following:

- The internal WCS TFTP server and the controller
- An external TFTP server and the controller

The use of the internal WCS TFTP server is probably the easiest and most straightforward method for most users. This option allows for the file to be loaded onto the TFTP server via one of the following two ways:

- Via the use of a TFTP client (the traditional approach)
- Using your client browser to transfer the file to the TFTP server home directory via HTTPS. With this method, a two-stage transfer is used between a directory on your local workstation and the WCS TFTP server to load the file onto the network device.

If you are updating multiple controllers in multiple download sessions throughout the day, it is more efficient to transfer the file from your desktop to a TFTP server only once and then specify a single stage transfer from the home directory of the WCS TFTP server for all subsequent controller downloads. (This is because when the source of a downloaded file is your local client workstation (“local machine”), the amount of traffic is increased twofold. The reason for this is because the software file is transferred twice: once between your desktop and WCS server using HTTPS, and then again between the controller and the WCS server using TFTP.)

The remainder of this section describes how to do this along with other options.

In the majority of centralized WLAN management implementations, the network traffic that results from the maintenance of controller software should not represent a major traffic component, especially with modern high-bandwidth campus LAN implementations. In larger implementations that may use slower or more congested WAN links between campus network and remote sites, it may be beneficial to transfer controller software files across the WAN to local TFTP servers during off-peak periods, especially if it is forecast that controller upgrades might need to be performed during peak traffic periods. By placing software files on TFTP servers that may be local to the network devices that require upgrading, transferring the files across the WAN during periods of peak traffic can be avoided.

In rarer cases of very large networks (such as those considered in Using WCS to Efficiently Deploy Your Wireless Network, page 8-92), additional considerations may be warranted. Given that controller operating software files are typically 25 MB in size, Cisco recommends that the performance impact of initiating multiple simultaneous controller operating software downloads during periods of peak network usage be more carefully examined. Multiple controller selections result in the initiation of multiple TFTP sessions (up to any limitation imposed by the TFTP server being used). The WCS administrator should keep this behavior in mind and limit the number of controllers selected, considering the underlying network topology, the bandwidth available, and other users on the network. As with other WCS displays, all controllers selected for software download must be present on one WCS display page when using the **Configure > Controllers** menu selection.

Perform the follow steps to download new software to WLAN controllers:

Step 1 Click **Configure > Controllers** to display the All Controllers page (shown in Figure 8-19). In large networks, you may find it helpful to use the “Search for Controllers” filtering feature in the left-hand margin to narrow the selection of displayed controllers by name, IP address, or network.

Figure 8-19 All Controllers WCS Page

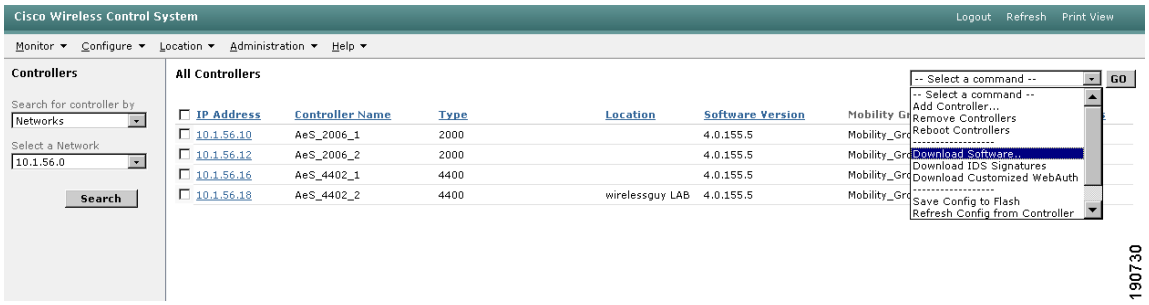

Step 2 The current version of controller software in each controller is listed under the column heading "Software Version". Enable the check box(es) to select the desired controller(s), choose **Download Software** from the command drop-down menu selector in the upper right-hand corner, and click **GO**. WCS displays the "Download Software to Controller" page.

Step 3 There are three choices with regard to transferring the software file to the WLAN controller(s) with the best choice being dependent on where the software file is resident:

- The software file is resident on your local computer—The file is resident on the client workstation that you are currently using to access WCS. In this case, you may use the two-stage process described previously to easily transfer the software file from your workstation to the WLAN controller.

  To do this, ensure that **Local Machine** is selected for the "File is Located on" option, as shown in Figure 8-20. Then click **Browse** to select the software file on your local computer. The file is transferred from your local machine to the internal WCS TFTP server, and then transferred to the controller.

Figure 8-20 Downloading Controller Software From the Local Machine

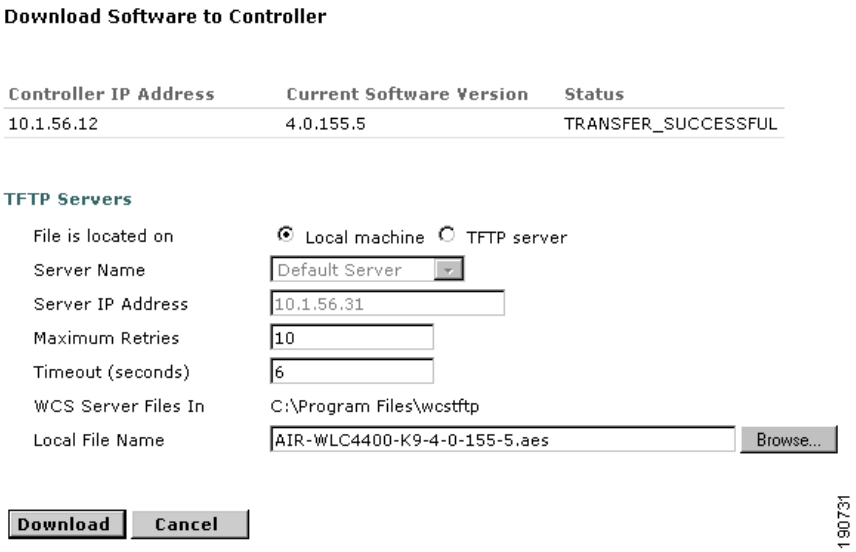

- The software file is already resident in the TFTP directory on WCS—This is the directory that you specified during the installation of WCS. Under this option, the software file is simply transferred from the WCS TFTP directory to the WLAN controller using TFTP only. To accomplish this ensure that:

- The **TFTP server** option is selected for the "File is Located on" option.
- "Default Server" is selected for the server name.
- The server IP address specified is the IP address of your WCS server. If it is not, it can be changed by modifying the template located at **Configure > Controller Templates > TFTP Server > Default Server**.
- The exact name of the file you want to load into the WLAN controller (such as *AIR-WLC4400-K9-4.0.155.5.aes*) is specified as shown in Figure 8-21. This filename needs to match the name of the file on the WCS TFTP server.

Figure 8-21 Downloading Controller Software from WCS TFTP Server

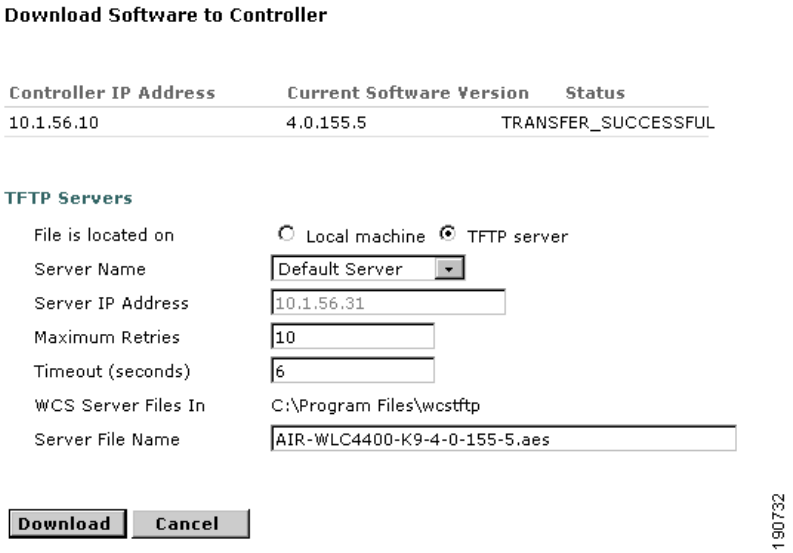

  - The use of this option is an efficient choice if you have already used option (3a) once before and the software file is now already resident on the WCS TFTP server, because it avoids re-transmitting the software file from the local machine to the WCS TFTP server unnecessarily.

- The software file is already resident in the TFTP directory of an external TFTP server—In this case, the file is transferred from the external TFTP server to the WLAN controller using TFTP.

Note Note that WCS cannot transfer a file from the local machine to an external TFTP server using HTTPS. Therefore, ensure that the file is already resident on the external TFTP server.

To accomplish this, ensure the following:

- The **TFTP server** option is selected for the "File is Located on" option.
- The external server of choice is selected from the drop-down menu for server name. If the external TFTP server you wish to use has not been defined on this WCS previously, it can be defined at this time by using "New" as the selection for the server name, typing in a name for this TFTP server definition and entering the server IP address.
- The exact name of the file you want to load into the WLAN controller (such as *AIR-WLC4400-K9-4.0.155.5.aes*) is specified as shown in Figure 8-22. This filename needs to match the name of the file on the WCS TFTP server.

Figure 8-22 Downloading Controller Software From an External TFTP Server

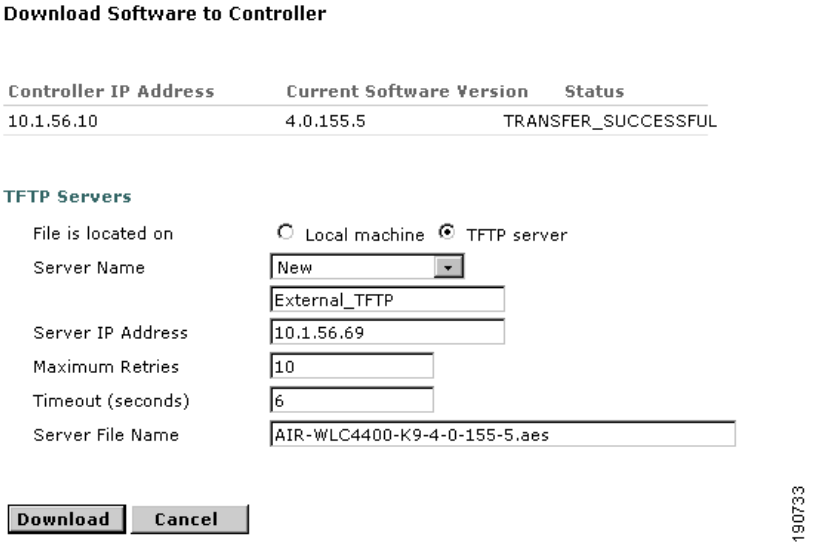

Step 4 After selecting the appropriate option in step 3, click **Download**. WCS downloads the software to the controller, and the controller initiates a process that ultimately results in the new software being written to nonvolatile (flash) memory. As WCS performs this function, it displays its progress in the Status field.

Step 5 After the download is complete, you need to save the current controller configuration (if desired) and reboot the controller for the new software to take effect. This can be easily performed by returning to the All Controllers screen shown in Figure 8-3, selecting the controller(s) you wish to reboot, selecting **Reboot Controllers** from the upper right-hand drop-down menu, and clicking **GO**.

Using the same basic steps outlined above, WCS also allows for web authentication bundles and intrusion detection (IDS) signatures to be downloaded to the controller as well in an analogous fashion. The commands to perform these functions can be accessed from the **Configure > Controllers > All Controllers** screen shown in Figure 8-19.

Further guidance concerning how to download controller operating software, web authentication bundles, and IDS signatures can be found in the *Cisco Wireless LAN Controller Configuration Guide, Release 4.0* and also in the WCS main menu bar under **Help > Online Help**.

Managing Location Server Software Level

As mentioned previously, WCS is the user interface to the location appliance and as such provides the control mechanism through which the location application and history databases on the location appliance are managed. All such software level management is performed for the location appliance from the **Location > Location Servers > Maintenance** screen shown in Figure 8-23. The Maintenance category shown provides several sub-category options for downloading new operating system software to the location appliance as well as performing a backup and restore of historical data on the appliance.

Figure 8-23 Location Server Maintenance

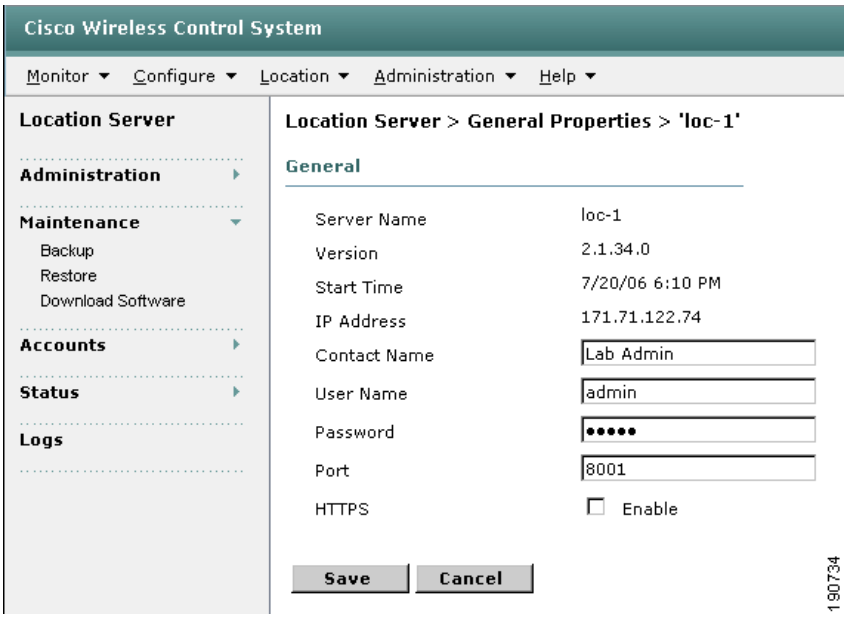

Complete step-by-step guidance about the updating of operating system software as well as how to perform appliance database backup and restore using WCS is available in the *Cisco Wireless Location Appliance—Configuration Guide*.

Ensuring Configuration Integrity

To ensure consistency, Cisco recommends that WCS be used whenever possible to configure and maintain the components of your Cisco Unified Wireless Network instead of direct CLI and GUI device access. As part of their configuration management functionality, many network management systems (including WCS) possess both an internal database structure where the last known configurations of network components are stored as well as the ability to query those components for their actual current configuration. Under normal operating circumstances, in an ideal environment where configuration access to network components is tightly controlled and only allowed from authorized network management stations, there should be little if any discrepancy between the actual configuration information contained in each network component and the configuration contained in the management system database.

In the real world, such discrepancies can and do arise. Whether from access by another group within the organization performing troubleshooting or from misconfiguration during hardware replacement, such situations may occur more often than is desirable in real-world deployments. To maintain integrity and value, an enterprise wireless network management system must possess a configuration management subsystem with which such discrepancies can be quickly identified and efficiently resolved.

When using the **Configuration > Controllers** function under the WCS main menu bar to manage WLAN controller configurations, the configuration object values that are displayed originate from the WCS internal database, not the controllers themselves.

Because WCS stores its representation of current WLAN controller configurations apart from the actual values present in the devices, the state of synchronization should occasionally be validated between the WCS databases and the actual managed device and if necessary, a re-synchronization should be initiated. WCS provides the network administrator with several tools to accomplish this. The subsections that follow describe these tools, which include the following:

- Configuration audit reporting
- Configuration synchronization
- WCS configuration refresh
- WLAN controller and access point configuration restoration

Configuration Audit Reporting

Configuration audit reports compare the complete current running configuration of a controller and its registered access points with the configuration stored in the WCS databases. Any exceptions are noted and brought to the attention of the network administrator via screen reports.

WCS offers both an on-demand as well as a periodically scheduled configuration audit reporting feature.

On-Demand Configuration Audit Reporting

To initiate an on-demand configuration audit report on behalf of a WLAN controller and its registered lightweight access points, follow these steps:

Step 1 Click **Configure > Controllers** to display the All Controllers page. In large networks, you may find it helpful to use the "Search for Controllers" feature in the left-hand margin to narrow the selection of displayed controllers by name, IP address, or network.

Step 2 Click on the hyperlink representing the IP address of the WLAN controller that you want to report. Note that a controller cannot be selected for this function by simply enabling the check box.

Step 3 Expand the **System** category selection in the left-hand column of the Controller Properties page. Click on the **Commands** subcategory, which brings up the Controller Commands page.

Step 4 Select **Audit Config** from the **Configuration Commands** drop-down selector and click on **GO.**

The result is a configuration audit report listing any discrepancies found, as shown in Figure 8-24.

Figure 8-24 On Demand Audit Report

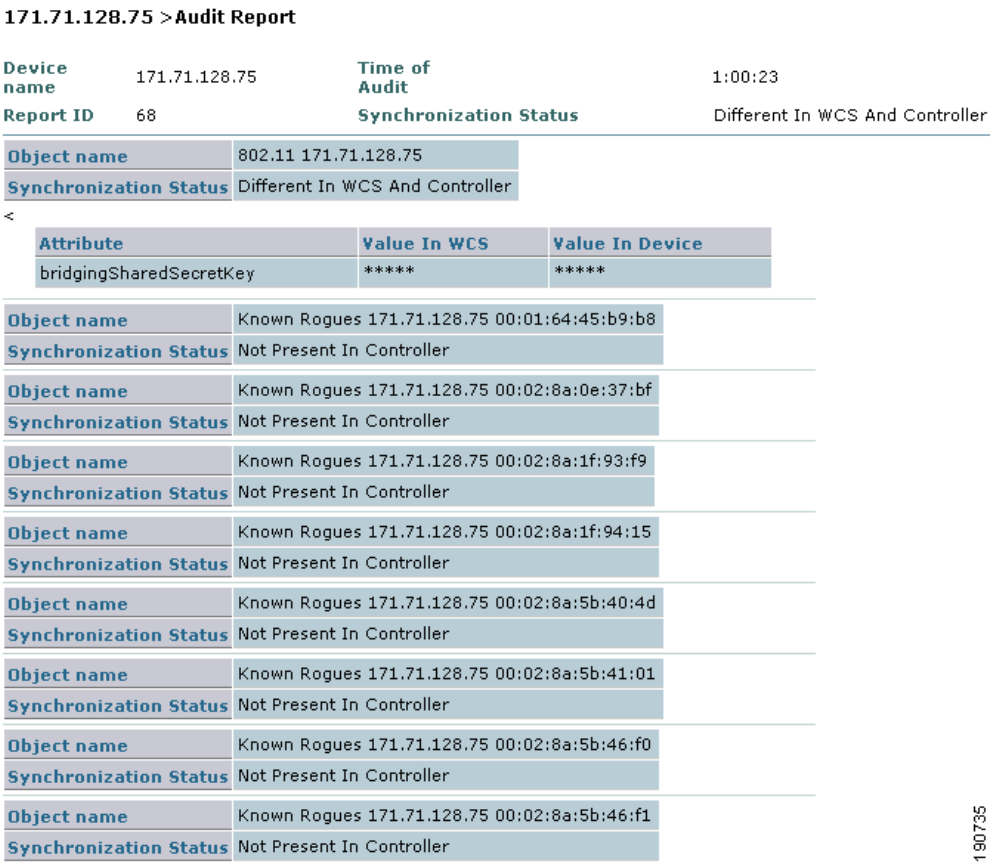

171.71.128.75 >Audit Report

Device name	171.71.128.75	Time of Audit	1:00:23
Report ID	68	Synchronization Status	Different In WCS And Controller

Object name	802.11 171.71.128.75
Synchronization Status	Different In WCS And Controller

<

Attribute	Value In WCS	Value In Device
bridgingSharedSecretKey	*****	*****

Object name	Known Rogues 171.71.128.75 00:01:64:45:b9:b8
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:0e:37:bf
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:1f:93:f9
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:1f:94:15
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:40:4d
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:41:01
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f0
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f1
Synchronization Status	Not Present In Controller

190735

Scheduled-Task Network Audit Reporting

WCS can also produce configuration audit reports automatically on a routine basis without user intervention. The output of this report is very similar to what has just been described. Known as the network audit report, this runs as a scheduled task and reports on discrepancies found between the configuration values in WCS databases and *all* WLAN controllers defined to WCS and their currently registered lightweight access points. Unlike the on-demand configuration report, the network audit report is non-selective and reports against all defined controllers that are SNMP reachable.

The network audit report can be configured to execute at a pre-defined time and with a pre-defined repetition interval. Alternatively, it can also be executed on a one-time “execute now” basis (which is equivalent in function to what was discussed in On-Demand Configuration Audit Reporting, page 8-33). The network audit report facilitates running unattended reports at times when network utilization is low and allows the viewing of report output to be deferred.

For further information on this scheduled task configuration audit report capability, see Network Audit, page 8-118.

Synchronizing WCS with Controller and Access Point Configurations

Synchronization in WCS is performed as a distinctly separate operation in relation to the identification of discrepancies. WCS offers several mechanisms through which configuration discrepancies between WCS, controllers, and lightweight access points can be resolved without requiring the operator to initiate

a manual configuration change. These options can be broken into two groups. The first of these is a selective synchronization option where the administrator may individually audit and synchronize select portions of controller and lightweight access point configurations. Access is via the same controller and lightweight access point configuration menus used to edit the configurations as described in Configuring WLAN Controllers, page 8-12 and Configuring Lightweight Access Points, page 8-16. The second group of options are non-selective one-way mechanisms that can either refresh the content of the WCS database from the controller configuration or restore the controller configuration from the information contained in the WCS databases.

Selective Synchronization

Selective synchronizations options are available on the same WCS menu panels that are used to specify parameters for controller and lightweight access point configuration. To initiate selective synchronization, access the screen of interest for the particular parameter category you want to audit by following the procedures already outlined in Configuring WLAN Controllers, page 8-12 and Configuring Lightweight Access Points, page 8-16. When you arrive at the parameter definition screen, you should notice that an "Audit" option is available alongside the option to save the configuration objects. An example is shown in Figure 8-25.

Figure 8-25 Selective Audit Option Example

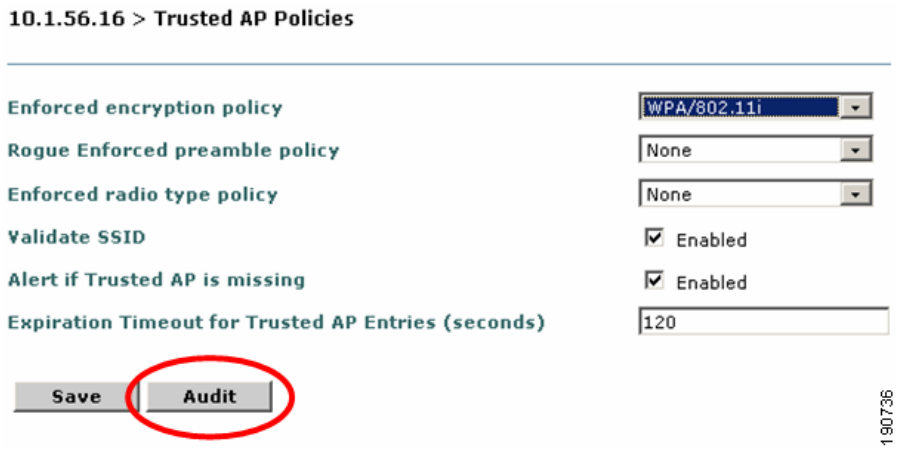

Selecting "Audit" in this case does not perform simple audit reporting, as was discussed in previous sections, but rather initiates the process of synchronizing the values of the displayed configuration parameters between the WCS database and the actual device running configuration. When a selective synchronization audit is performed, WCS queries the WLC and requests the audit parameter values contained in the WLC for the specified Management Information Base (MIB) objects. WLC responds with the current values for the audit parameters. When the responses are received, WCS compares the values contained in its databases to the values received from the device.

If the results of the comparison indicate that the device is in synchronization with WCS, WCS indicates that no differences exist. However, if there are any discrepancies found during the comparison, WCS presents the operator with a screen similar to that shown in Figure 8-26.

Figure 8-26 Auditing a Configuration Category

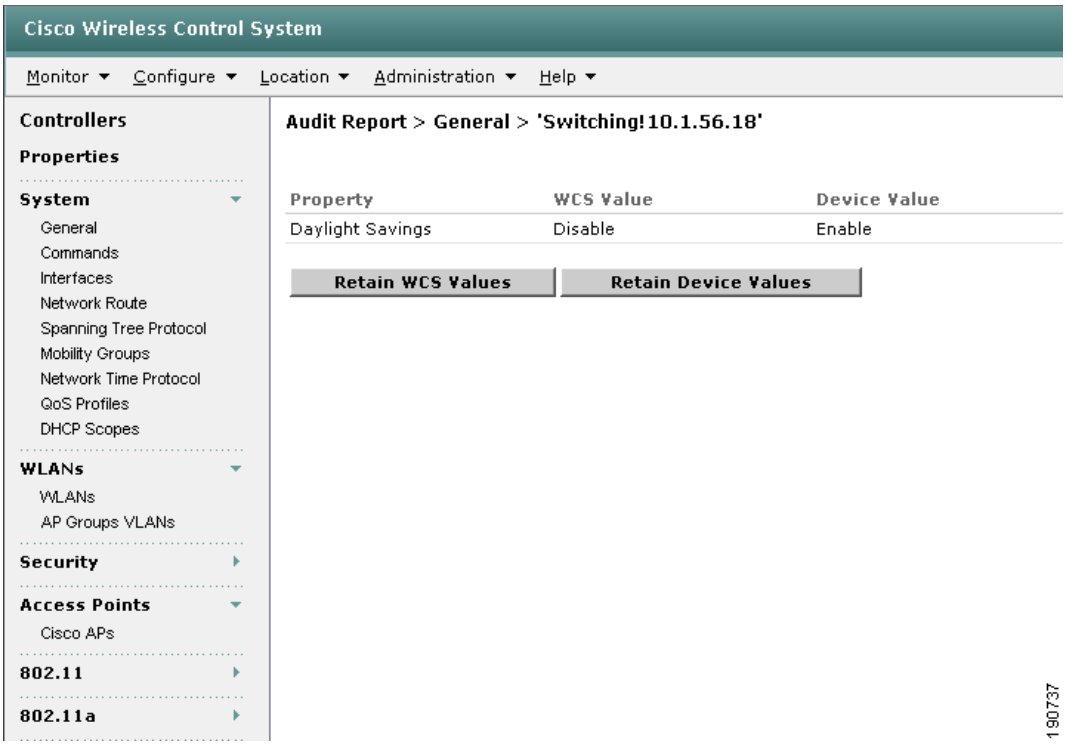

Clicking on one of the two options in Figure 8-26 causes one of the following to occur:

- Retain Device Values—The information in WCS that conflicts with the information shown in the device is overwritten with the information specified in the device.
- Retain WCS Values—The information in the device that conflicts with the information shown in WCS is overwritten with the information specified in WCS.

In either case, WCS presents confirmation of the selection and the action performed as a result.

In this way, WCS allows the operator to review as little or as much of the device configuration as desired, and synchronize only selected parts of the configuration.

This same procedure can be used to perform selective synchronization of lightweight access point configurations as well. See Configuring Lightweight Access Points, page 8-16 for information about the configuration of lightweight access points.

Non-Selective Synchronization

In some cases, you may wish to perform synchronization between WCS and its managed resources on a grander scale than that which is available via selective synchronization. WCS offers the capability of performing one-way non-selective synchronizations of the WCS database with the entire running configuration contained within the controller (or vice-versa).

- Refresh Config from Controller—When the **Refresh Config from Controller** feature is selected, a one-way (WLC -> WCS) synchronization of all configuration objects for the selected WLAN controller is performed. A one-way synchronization in this case implies that all configuration information pertaining to the controller on WCS is overwritten with the running configuration of the WLAN controller. This feature is useful in correcting a situation where the WCS database has become out of sync with the configuration contained in the controller in multiple configuration

categories. This can result, for example, if changes are made to the WLAN controller configuration in WCS but because of a communication or other errors in the controller, the changes were not completely applied. WCS and the WLC operate perform such updates in unison and with close monitoring of error status to preclude the occurrence of such events. But if this type of situation should occur, **Refresh Config from Controller** provides a simple way to completely re-synchronize the controller information contained within WCS to the current running configuration of the device.

Refresh Config from Controller can be performed against either a single WLAN controller or multiple WLAN controllers simultaneously. To perform the refresh, follow these steps:

Step 1 Click **Configure > Controllers** to display the All Controllers page. In large networks, you may find it helpful to use the "Search for Controllers" filter in the left-hand margin to narrow the selection of displayed controllers by name, IP address, or network.

Step 2 Enable the check box(es) for each WLAN controller(s) whose configuration(s) you want to restore to the WCS database. Keep in mind that for each controller selected, WCS sends a series of SNMP PDUs to retrieve the running configuration of each controller.

Step 3 From the command drop-down menu selector in the right-hand upper corner of the screen, choose **Refresh Config from Controller** and click **GO.**

The page shown in Figure 8-27 is presented.

Figure 8-27 Refresh Config from Controller Conflict Resolution

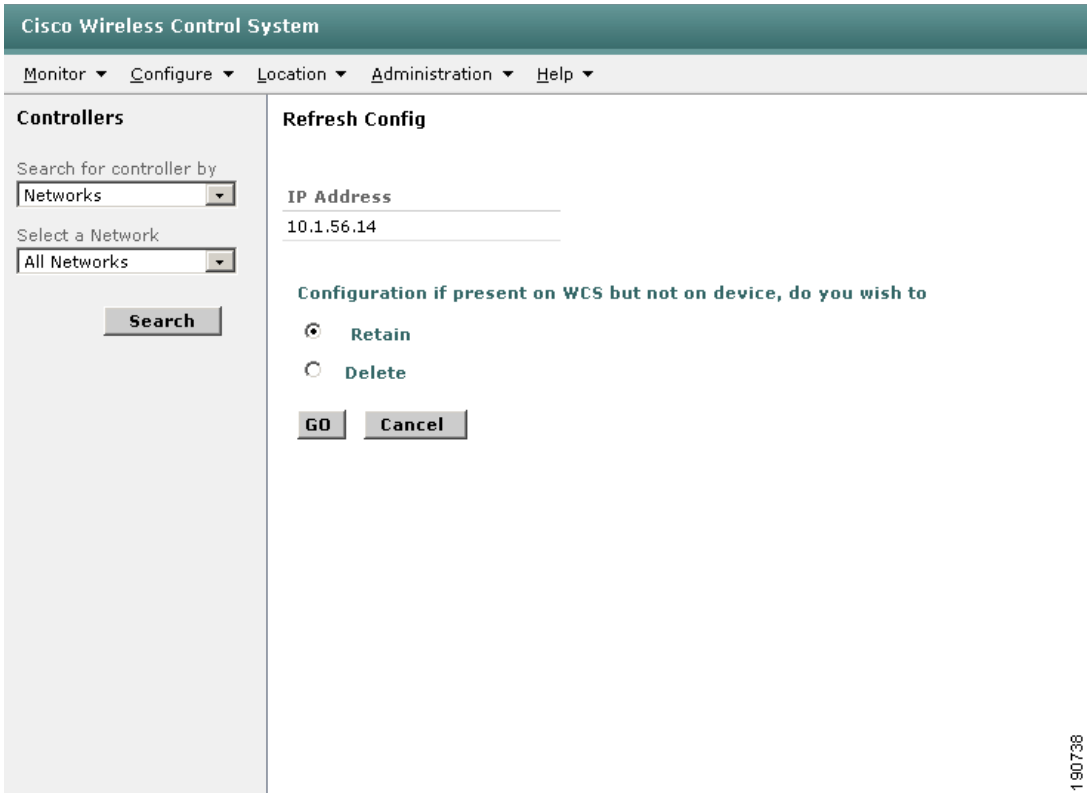

This information displayed concerns itself with what to do in the event that a configuration object exists in the WCS database for the controller but does not exist in the running configuration of the controller.

Step 4 Do one of the following:

a. Select **Retain** if you want the value found in the WCS database to prevail.

b. Click **Delete** if you want to remove the existing value found in the WCS database and replace it with the value found in the controller.

Step 5 Click **GO** after you have made your selection.

- Restore Config to Controller—This provides for a one-way, non-selective synchronization of all applicable controller configuration objects from the WCS database to the running configuration of the selected WLAN controller is performed. All information contained in the controller is overwritten with the information contained in the WCS databases.

  Unlike **Refresh Config from Controller**, **Restore Config to Controller** can only be applied on an individual basis against specific controllers.

  To refresh the running configuration of a WLAN controller from the WCS database using **Restore Config to Controller,** perform the following procedure:

Step 1 Click **Configure > Controllers** to display the **All Controllers** page.

In large networks, you may find it helpful to use the "Search for Controllers" feature in the left-hand margin column to narrow the selection of displayed controllers by name, IP address, or network.

Step 2 Click the hyperlink representing the IP address of the WLAN controller that you want to configure.

Step 3 Expand the **System** category selection in the left-hand column of the **Controller Properties** page, as shown in Figure 8-28. Click the **Commands** subcategory, which brings up the **Controller Commands** menu.

Figure 8-28 Controller Commands

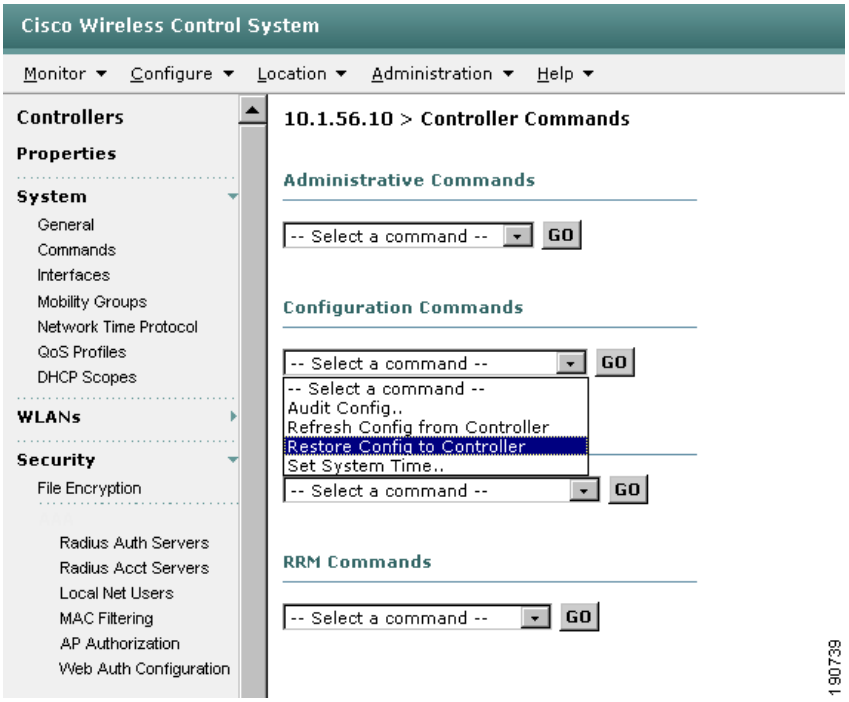

Step 4 Select **Restore Config to Controller** from the **Configuration Commands** drop-down selector and click **GO**.

Be careful not to inadvertently select **Refresh Config from Controller** instead of **Restore Config to Controller** because both options appear in the drop-down menu selector.

Step 5 Confirm the action by clicking **OK** on the confirmation screen.

Controller Configuration Archival

The ability to identify discrepancies between the contents of WCS and the actual configuration of network devices, coupled with two powerful mechanisms allowing for re-synchronization is usually sufficient to recover from most accidental or unintentional out-of-sync situations. In some cases, however, the out-of-sync situation can become somewhat aggravated because of the passage of time or the inadvertent operator acceptance of controller configuration changes into the WCS database that are later found not to have been valid. To address these situations and others, WCS also provides the ability for archival and restoration from external configuration archive files that are independent of the WCS database itself.

Note WCS also provides the ability to backup and restore the WCS databases themselves to protect against the rare occurrence of WCS database failure or corruption. This section concerns itself more with isolated cases of single controller configuration corruption, and not cases of widespread WCS database corruption.

Regular archiving of device configuration is a best practice because it allows the recovery of lost configuration, and it also provides an audit trail of when changes occurred in device configuration. This section describes how WCS provides for controller configurations to be archived to individually named and time-stamped files on a designated TFTP server. The archival process can be initiated either on-demand or via a scheduled task.

Archiving of Individual Controller Configuration Files

WCS provides the ability to manually archive the configuration of a WLAN controller to a file on a designated TFTP server via the following process.

Step 1 Click **Configure > Controllers** to display the All Controllers page. In large networks you may find it helpful to use the "Search for Controllers" feature in the left-hand margin to narrow the selection of displayed controllers by name, IP address, or network.

Step 2 Click the hyperlink representing the IP address of the WLAN controller for which you want to archive the configuration. Note that for this function the controller cannot be selected by simply enabling the check box.

Step 3 Expand the **System** category selection in the left-hand column of the Controller Properties page. Click on the **Commands** subcategory, this brings up the Controller Commands page shown in Figure 8-28.

Step 4 Select **Upload File From Controller** from the **Upload/Download Commands** drop-down selector and click on **GO**. This displays the screen shown in Figure 8-29.

Figure 8-29 Manually Archiving Controller Configurations, Logs, and Signature Files

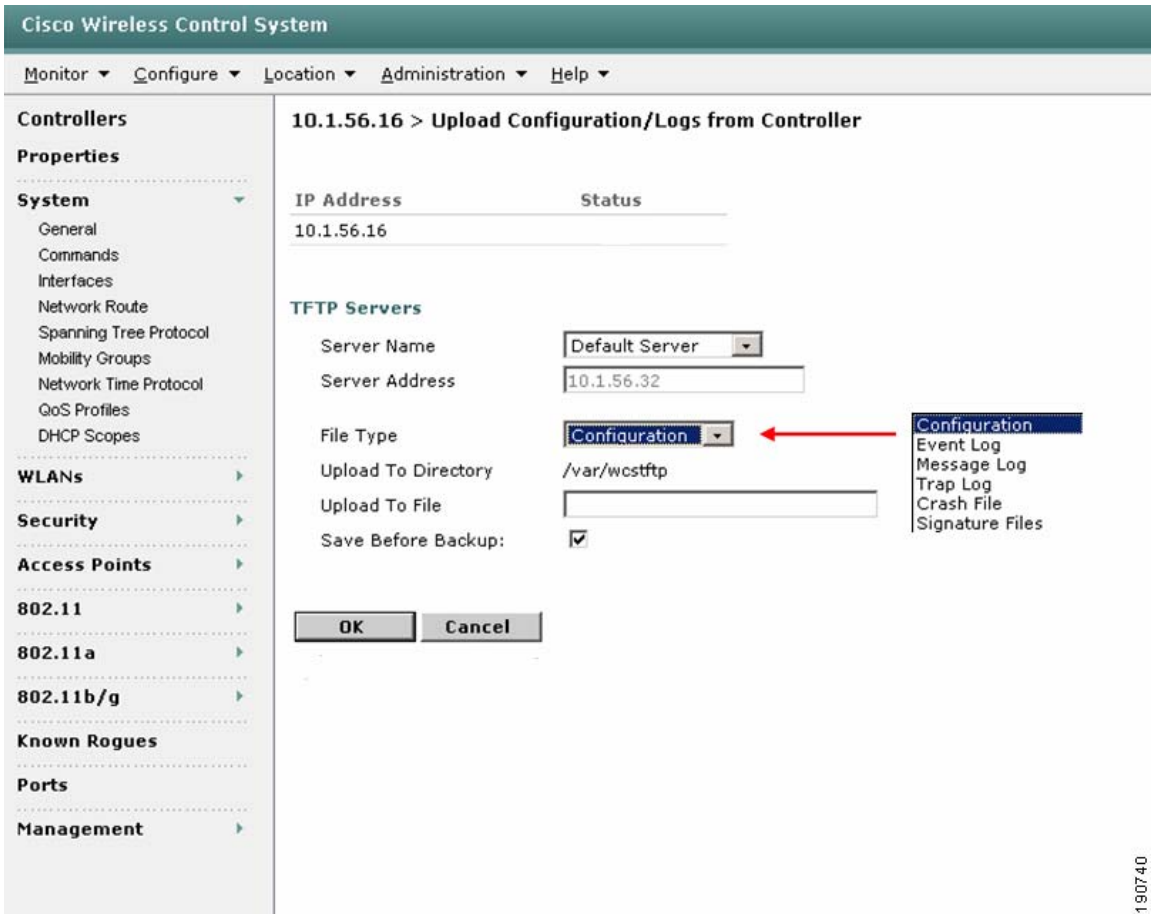

Step 5 Select a server that has been already configured from the Server Name drop-down selector, or select "New" to define a new TFTP server and enter the IP address. For file type, select **Configuration**. Note that this same mechanism can be used to archive other controller files such as IDS signature files and logs.

Step 6 Enter the file name that you want the archive saved as on the TFTP server.

Step 7 You may wish to enable the **Save Before Backup** check box. When this check box is enabled, it indicates that the running configuration of this controller should be saved to the internal nonvolatile (flash) memory of the controller before the archiving begins. Because the scheduled task archives only the *saved configuration of the controller and not the currently running configuration*, enabling this check box ensures that any recent unsaved changes are included in the archive.

Step 8 Click **OK**.

You may see a warning about enabling file encryption on this page. AES file encryption can be configured using **Configure > Controllers > Security > File Encryption**. File encryption is highly recommended when archiving controller configurations over WANs and other public communications facilities.

Automatic Archival of Controller Configurations

WCS also provides a automated routine (known as the **Configuration Backup** scheduled task) that automatically archives the configuration of each reachable controller that has been defined to WCS.

Note For further information, see section 13.6 CSCsd54800—Config Backup Scheduled Task Fails if One Controller is Unreachable.

The Configuration Backup scheduled task can be configured to run at a pre-defined time of day and with a pre-defined repetition interval. It can also be submitted for execution on a on-demand basis.

See Configuration Backup, page 8-117 for further details on the **Configuration Backup** scheduled task.

Restoring Controller Configuration Archives

WCS provides the ability to manually restore an individual configuration archive to a controller via **Configure > Controllers > Commands > Upload/Download Commands > Download Config**. Note that controller configuration archives can only be restored on an individual, one-at-a-time basis.

Configuring WCS Campus, Building, Outdoor, and Floor Maps

Cisco WCS allows for the addition of *maps* to its internal database that can then be used to assist in the visualization of client, asset tag, and rogue location as well as estimated coverage during the monitoring of your wireless LAN. Adding maps to the Cisco WCS database enables you to view your managed system on realistic outdoor, campus, building, and floor plans that you have defined that allow more meaning to be imparted to the viewer. Maps can originate from actual floor plans that are imported into WCS using .PNG, .JPEG, .JPG, or .GIF graphic file formats (AutoCAD .DXF file formats are not supported at this time). After they are imported and sized, RF characteristics can be added to various building components to increase coverage prediction and design accuracy.

Maps are usually added in campus, building, and floor sequence; however, the existence of a campus map is not mandatory (buildings can be freestanding and not part of a campus in smaller designs). Floor maps cannot exist independently of building maps. Outdoor areas are typically associated with campus maps and do not exist independently.

The WCS *map editor* can be used to define, draw, and enhance floor plan information. The map editor enables the creation of obstacles that can be taken into consideration when computing RF prediction heat maps for access points. (Although supported, Cisco recommends the use of the map editor to draw walls and other obstacles rather than importing .FPE files from the legacy floor plan editor.) You can also add coverage areas that are used by the location appliance to locate clients and 802.11 active RFID tags and provide alarm notifications on their movement.

WCS *planning mode* is a feature that enables you to calculate the number of access points required to cover an area by placing fictitious access points on a map and allowing you to view the coverage area that they would yield. Based on the throughput specified for each protocol (802.11a or 802.11b/g), planning mode calculates the total number of access points required to provide optimum coverage in your network.

Detailed step-by-step guidance on how to define, add, and edit campus, building, and outdoor floor maps can be found in the *Cisco Wireless Control System Configuration Guide, Release 4.0*.

Configuring WCS to Manage the Cisco Wireless Location Appliance

The location appliance enhances the high-accuracy location capabilities that are integrated into location-enabled WCS servers. The location appliance enhances WCS location capabilities by computing, collecting, and storing historical location data and allowing WCS to display location data for multiple tracked devices at a time. In addition, the location appliance handles the dispatch of location notifications and provides the SOAP/XML API to which third-party applications can interface to the location information stored within the location appliance databases. Before use, the location appliance must be configured and defined to the WCS server that has been licensed for location services.

After initial IP parameter settings as described in the *Cisco Wireless Location Appliance—Installation Guide* and the *Cisco Wireless Location Appliance—Configuration Guide*, all configuration of the location appliance is performed from WCS using the menus and submenus located under the **Location** tab.

Complete guidance in configuring and managing the location appliance via the menus located under the WCS **Location** tab can be found in the *Cisco Wireless Location Appliance—Configuration Guide*. This includes step-by-step configuration instructions on the following topics:

- Adding and deleting location servers
- Synchronizing Cisco WCS and location servers
- Editing location server properties
- Managing location server users and groups
- Configuring location event notifications
- Monitoring location servers
- Performing location server maintenance

In addition, extensive coverage of location-based services and positioning technologies, location-aware design, deployment best practices and RFID tag technology are available in the following documents:

- Wi-Fi Location Based Services: Design and Deployment Considerations—This document can be found at http://www.cisco.com
- Cisco Wireless Location Appliance: Deployment Guide—http://www.cisco.com/en/US/products/ps6386/prod_technical_reference09186a008059ce31.html

Using WCS to Monitor Your Wireless Network

WCS facilitates the monitoring of device status within its management domain via several avenues. Whether looking at the omnipresent alarm counters that appear in the lower left-hand corner of every WCS screen or the detailed status information available under the **Monitor** tab, information about the current status of your enterprise network is presented to the operator in a clear, efficient, and visually attractive manner. Access to this information and more starts with a simple mouse click on the **Monitor** tab on the main menu bar (or the *Alt-M* keyboard shortcut).

The following sections briefly describe the information available to you under each of the selections in the **Monitor** menu tree. Keep in mind that as seen in the **Configuration** menus discussed previously, the majority of device status information in the WCS databases is accessible via multiple paths in the GUI. For example, although information about controller and lightweight access point status is readily available by clicking on **Monitor > Devices > Controllers** or **Monitor > Devices > Access Points** respectively, much of the same information is accessible via strategically located hyperlinks on many other **Monitor > Device** submenus as well.

WCS also provides a northbound notification system that can dispatch e-mails to specified destinations (pager, cell phone, PDA, laptop, and so on) whenever certain types of alarms occur.

The following sections describe the monitoring capabilities of WCS and how these can be put to use in monitoring the enterprise wireless network.

Network Summary

WCS typically presents a summary page at the top of the Monitoring menu tree called the *Network Summary***.** Network Summary is presented immediately after clicking on the **Monitoring** main menu selection (or using the keyboard shortcut Alt-M). As shown in Figure 8-30, Network Summary shows what is currently taking place in your wireless network that may warrant your immediate attention.

Figure 8-30 Network Summary

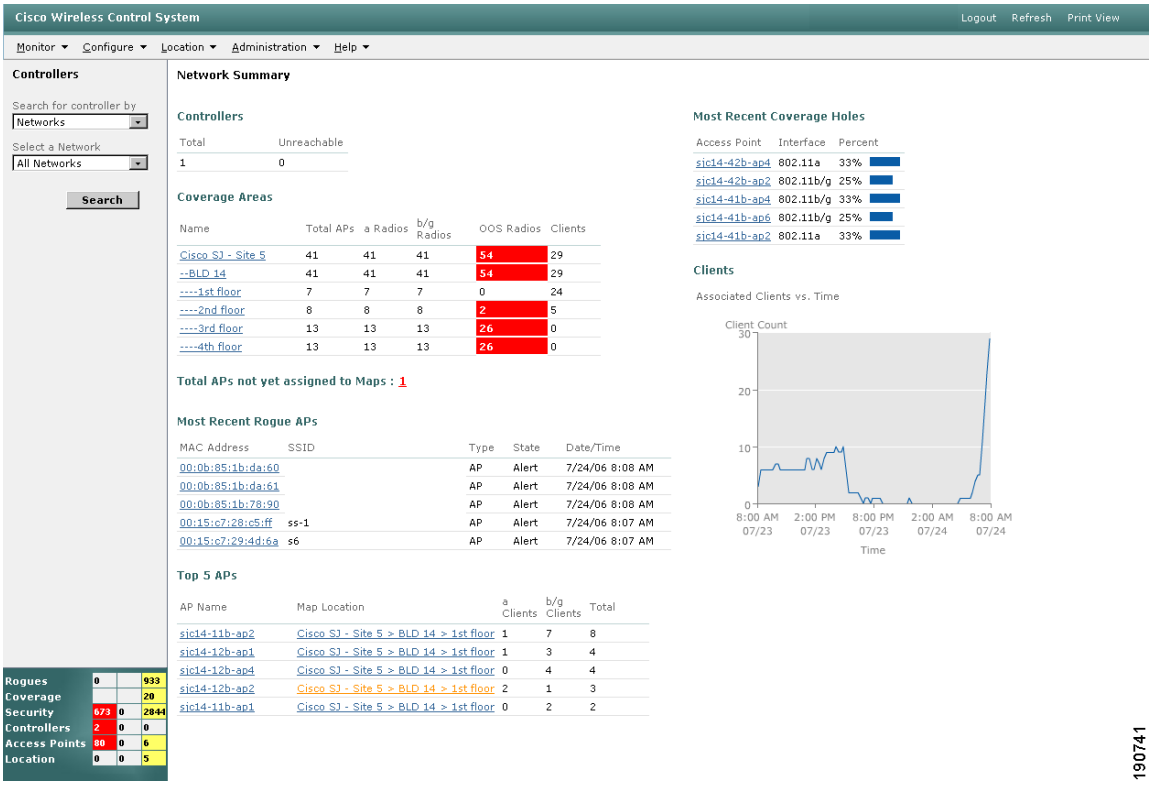

Along with the information that is presented via the alarm summary panel in the lower left-hand corner of the screen, you are shown a cross-sectional view into the status of several key areas that may affect not only the operational and performance characteristics of your wireless LAN but its overall security as well.

The following can be learned from the Network Summary panel:

- Controllers—Total number of WLAN controllers defined to WCS and the number of controllers that WCS has determined to be unreachable (that is, time-outs to SNMP queries from WCS). For more details on precisely which controllers are unreachable, click **Monitor > Devices > Controllers**.

- Coverage Areas—Up to ten "coverage areas" that have been defined to WCS along with the total number of lightweight access points, access point radios, and clients found. Any access point radios that have been administratively disabled or down for other reasons are listed in the Out-Of-Service (OOS) column.

Note The use of "coverage areas" in the context of the Network Summary page should not be confused with the coverage areas that are defined in the location appliance.

The coverage areas listing may consist of a combination of campuses, buildings, floors, or outdoor areas. If WCS determines that ten or more campus or standalone building maps have been defined, it displays a "View All Maps" hyperlink that enables you to jump to the **Monitor > Maps** page where the entire list of defined maps can be seen.

Clicking on any of the hyperlinks shown in this area allows you to move to the associated map screen where you are able to obtain detailed information on which lightweight access points in that area are experiencing difficulties and what those difficulties may be. In some cases, a hyperlink may appear indicating the total number of lightweight access points that have not been assigned to any maps.

- Most Recent Rogue Access Points—This area lists information concerning the five most recently detected rogue access point alarms and provides hyperlink access to the **Alarm > Rogue AP** page. The **Alarm > Rogue AP** page lists more detail about the detected rogue access point such as its location, event history, and any rogue clients that may be associated with it.
- Top Five Access Points—The current list of the top five lightweight access points ranked by the total number of client associations. From this list, an AP Name hyperlink provides access to the associated **Monitor > Devices > Access Points** panel for each lightweight access point. The location hyperlink takes you to the floor level map where the lightweight access point has been placed.
- Most Recent Coverage Holes—The names of the lightweight access points that have generated the five most recent "coverage hole" alarms. A "coverage hole" is an alarm situation triggered because of the crossing of a minimum signal coverage threshold by a client. When an alarm is triggered, it usually indicates that a client has entered an area where the minimum signal detected by the client from any of the lightweight access points servicing that area is below pre-determined threshold levels. When this is communicated to WCS via SNMP, it normally results in the generation of a coverage hole alarm. Clicking on any one of the AP name hyperlinks found in this area takes you directly to the detailed coverage hole alarms page that displays the current alarm status.

  Keep in mind that although a coverage hole alarm may have already been cleared, the cleared alarm is still listed in this area unless superseded by other coverage hole alarms.
- Associated Clients versus Time—This graph shows the total number of associated clients across the WCS management domain. You can perform a mouse-over of various points on the graph that display additional information in a pop-up bubble about the number of users associated and the date/time that the sample was taken. There are no hyperlinks associated with the points on the graph.
- Critical, Major, and Minor Alarms—Although discussed here, this multi-colored rectangular area is present in the lower portion of the left-hand margin of virtually every WCS display. It contains a tabulation of the number of critical (red), major (orange), and minor (yellow) alarms for each resource category (rogues, coverage, security, controllers, access points, and location). If all alarms of a given type are cleared in a particular category, the count reflects zero and the color changes to white. This feature has been found to be very useful in every network management because it constantly keeps the WCS administrator abreast of all alarm counts while consuming minimal screen real estate.

Monitoring Maps

As discussed in the previous section, WCS allows for visual status of devices to be displayed on campus, building, floor, and outdoor maps. The following examines how **Monitor > Maps** allows the WCS user to fully take advantage of these capabilities in managing the enterprise wireless LAN.

The main page under **Monitor > Maps** (shown in Figure 8-31) contains similar information to what was displayed in the Coverage Areas portion of the Network Summary screen.

Figure 8-31 Monitor > Maps

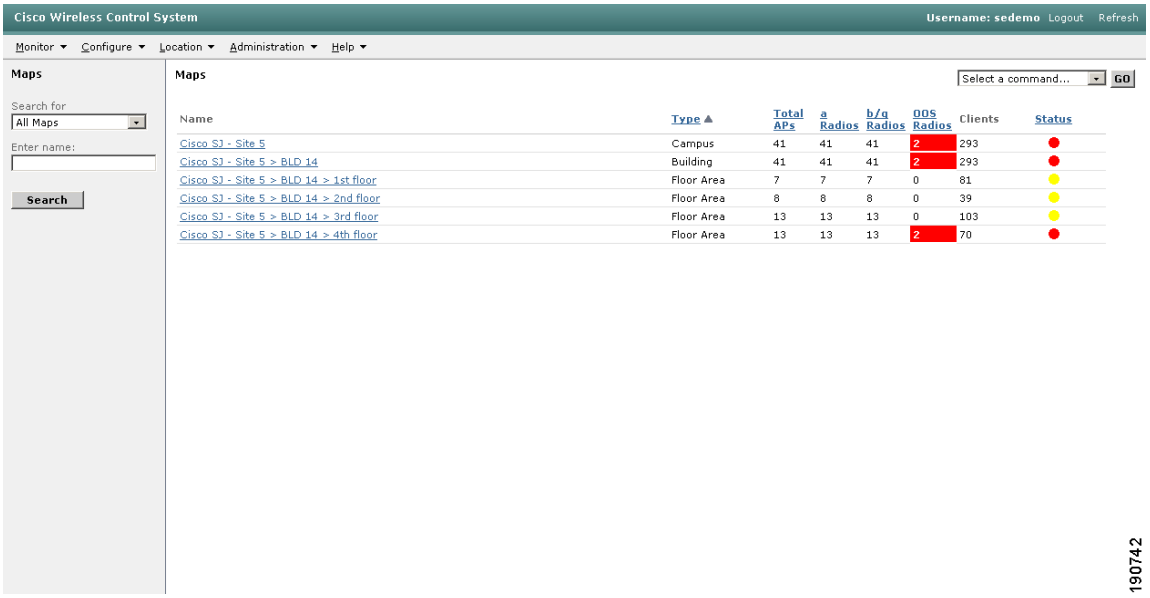

However, here you see the entire listing of all known maps instead of simply a summary excerpt. In large networks comprising more than a single campus with each in turn made up of several buildings with multiple floors, the viewer may find it easier to manage the contents of the display by using the "Search For" selector in the left-hand margin. The display content can also be sorted as per the viewers preference (default sort order is by type). This sort order can be changed by clicking on any of the column headings, which sort the display by the column values in either ascending or descending order.

Figure 8-31 shows that the visual alarm status color indication is provided for the indicated resources. Note that the alarm status of any map in the hierarchy depends not only upon the status of the resources on that map but also on the status of resources located on any lower level maps as well. Thus, if a resource on a floor map generates an event triggering a critical alarm, the campus and building maps as well as the floor map all show a red critical alarm status. In this manner, the presence of an alarm attracts the attention of the WCS administrator no matter what level of the campus hierarchy they happen to browse to in pursuit of the matter.

Clicking on any of the map name hyperlinks results in campus, building, floor, or outdoor area maps being displayed. Floor maps (shown in Figure 8-32) contain icons representing each lightweight access point. The color of these access point icons vary to reflect the state of alarm that has been triggered by events occurring at the lightweight access point.

Figure 8-32 Graphical Floor Map Showing AP Locations and Heat Maps

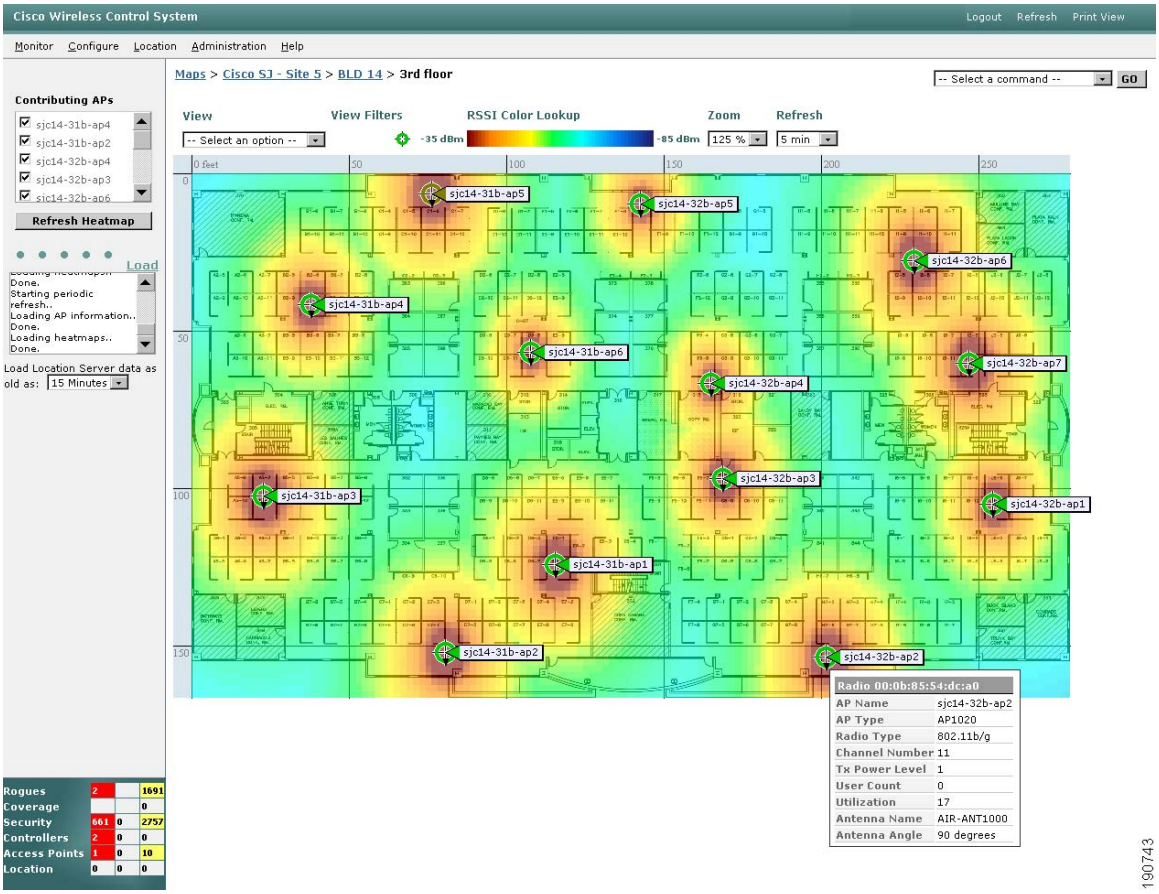

Floor maps can be viewed at various zoom settings (50–800 percent and full screen) with a screen refresh timer that can be configured from 5 seconds to 15 minutes. The following bullets outline just some of the information that WCS provides graphically using floor maps as shown in Figure 8-32:

- RF coverage predictions (heatmaps) based on current access point settings (with the ability to specify participating access points)
- Lightweight access point locations with antenna orientation and alarm status, and the ability to:
  - Filter by radio type (protocol of 802.11a, 802.11bg or both)
  - Display an icon label for each lightweight access point containing one of the following: channels, TX power, coverage holes, MAC address, AP name, controller IP, utilization percentage, profiles (load, noise, interference, and coverage), or user count
  - Limit heatmaps using an RSSI cutoff (not to be confused with **Location > Location Servers > Location Parameters > RSSI Cutoff**) of between -85 dBm to -60 dBm. This allows you to easily predict where the minimum acceptable RSSI for a particular wireless device (such as an 802.11 wireless phone) likely resides.

Using a version of WCS licensed for location-based services along with a location appliance adds graphical capabilities such as the following:

- Display location coverage areas and coverage markers
- Display the location of multiple clients simultaneously with the ability to:

  - Display as an icon label one of the following: IP address, username, MAC address, asset name, asset group, asset category, or controller IP address
  - Filter the displayed clients by IP address, user name, MAC address, asset name, asset group, asset category, SSID, radio type (protocol), and controller IP address
- Display of 802.11 asset tag locations with the ability to:
  - Display as an icon label one of the following: MAC address, asset name, asset group, and asset category
  - Filter the displayed asset tags by MAC address, asset name, asset group, asset category, and controller IP address
- Display of multiple rogue access point locations with the ability to filter by:
  - MAC address, on-network status (yes/no/either) or state (alert, known, acknowledged, contained, threat, or known contained)
- Display of multiple rogue access point clients with the ability to filter by:
  - Associated rogue access point MAC address
  - State (alert, contained, or threat)

Figure 8-32 also indicates a drop-down menu selector in the left-hand margin that allows all "live" location data retrieved from the location appliance to be filtered by age. In Figure 8-32, this is defaulted to 15 minutes but it can be set from 2 minutes to as long as 24 hours.

This is only a brief overview of the many capabilities available under WCS and its graphical monitoring facilities. Additional information pertaining to the entire range of information viewable under **Monitor > Maps** can be found in the *Cisco Wireless Control System Configuration Guide, Release 4.0*. Additional information can also be found in the WCS menu bar under **Help > Online Help**.

Information specific to the configuration of WCS floor maps for displaying the location of clients, asset tags, rogue access points, and rogue access point clients when using WCS with the Wireless Location Appliance can be found in the following documents:

- Wi-Fi Location Based Services: Design and Deployment Considerations—This document is available at http://www.cisco.com
- Cisco Wireless Location Appliance: Deployment Guide—http://www.cisco.com/en/US/products/ps6386/prod_technical_reference09186a008059ce31.html

Monitoring Devices

Thus far, this guide has examined the monitoring capabilities of WCS based primarily based on network geography. This section describes the capabilities available that are indexed instead by device category. This proves useful when you are primarily interested in seeing current alarm status for all devices comprising a resource category, regardless of where they may be located within the management domain (although you can filter the listings on location as well).

Clicking on **Monitor > Devices** on the main WCS menu bar shows that device status is grouped into four main device categories: Controllers, Access Points, Clients, and Tags. The subsections that follow take a brief look at each of these.

Monitoring WLAN Controllers

Clicking on **Monitor > Devices > Controllers** displays the page shown in Figure 8-33.

Figure 8-33 Monitor > Devices > Controllers

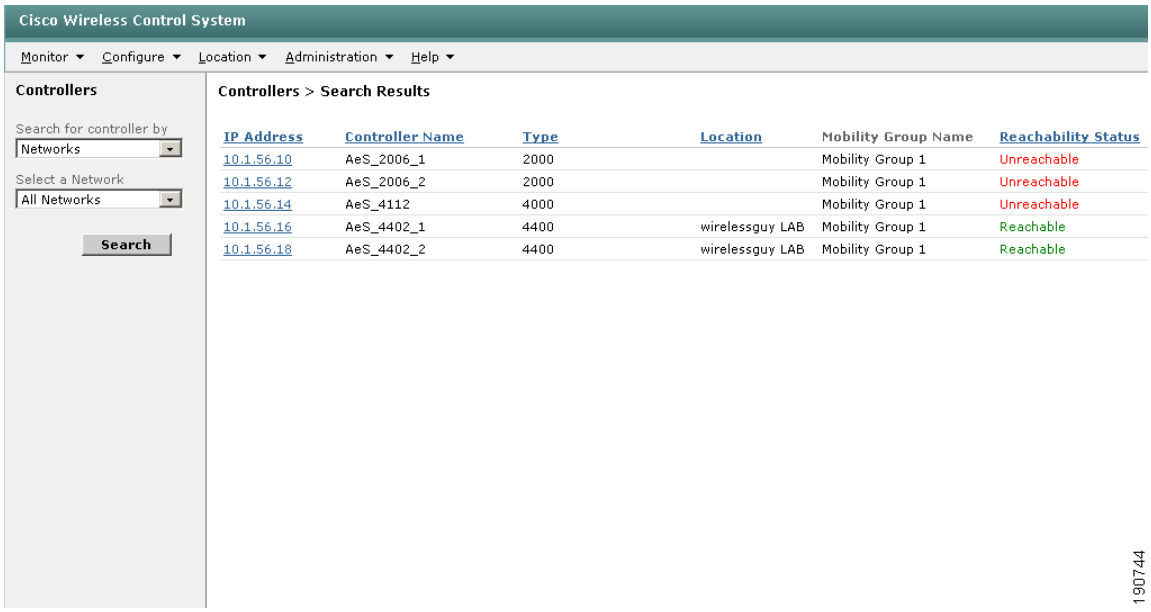

As seen previously in **Configure > Controllers**, Figure 8-33 indicates whether all controllers managed by this WCS indicate are SNMP reachable. A controller becoming unreachable is an event that triggers a critical alarm, so in this monitor screen it is shown in red. Clicking on the hyperlink for any controller IP address leads the WCS user to the **Controller > Summary** page, as shown in Figure 8-34. The **Controller > Summary** page is the main launching point from which you can drill down into more in-depth status information about a particular controller within the WCS management domain.

Figure 8-34 illustrates and example of a controller summary display for a Cisco WLAN controller model 4400 with 100 lightweight access point capacity.

Figure 8-34 Monitor > Devices > Controller Summary for 4400 Series WLAN Controller

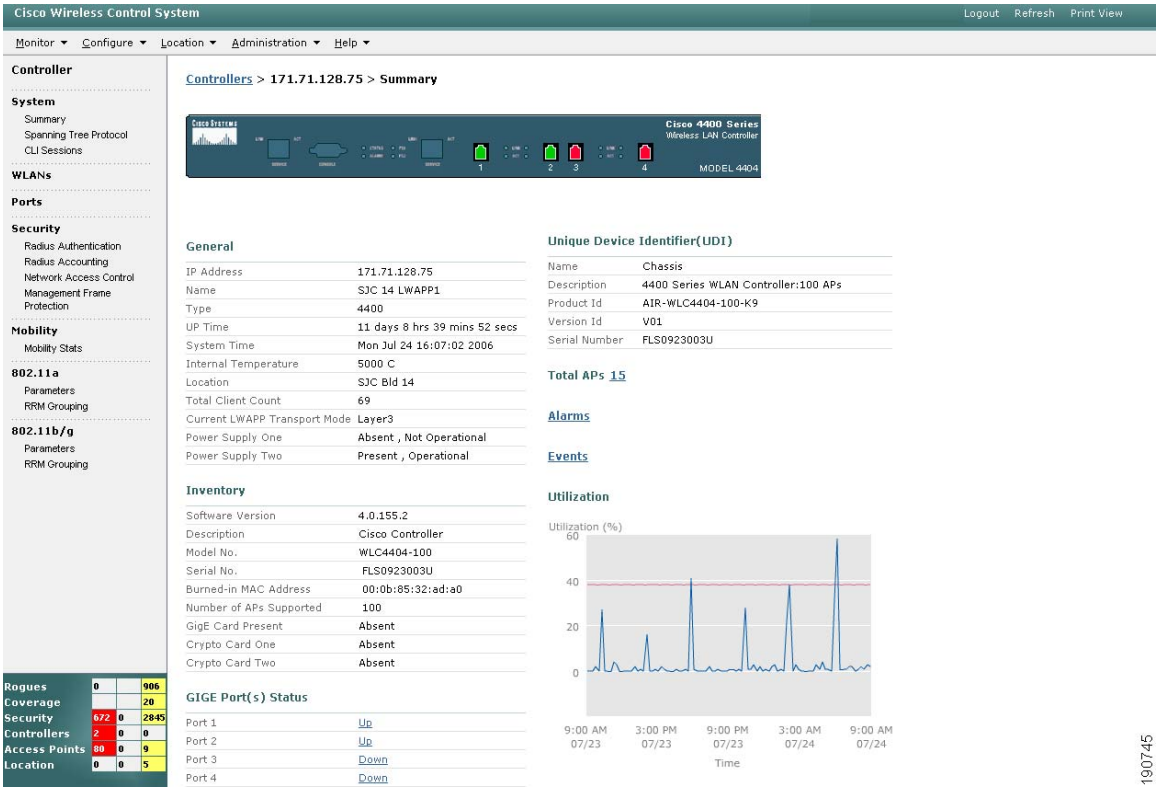

The graphic that appears varies depending on the model of the controller hardware (4400, 2000, Cat6500/WiSM, ISR/WLCM, and so on). For example, the Cisco Catalyst 6500 Wireless Solutions Module (WiSM) and its two onboard controllers can be visually represented as shown in Figure 8-35.

Figure 8-35 Monitor > Controller > Summary for Catalyst WiSM

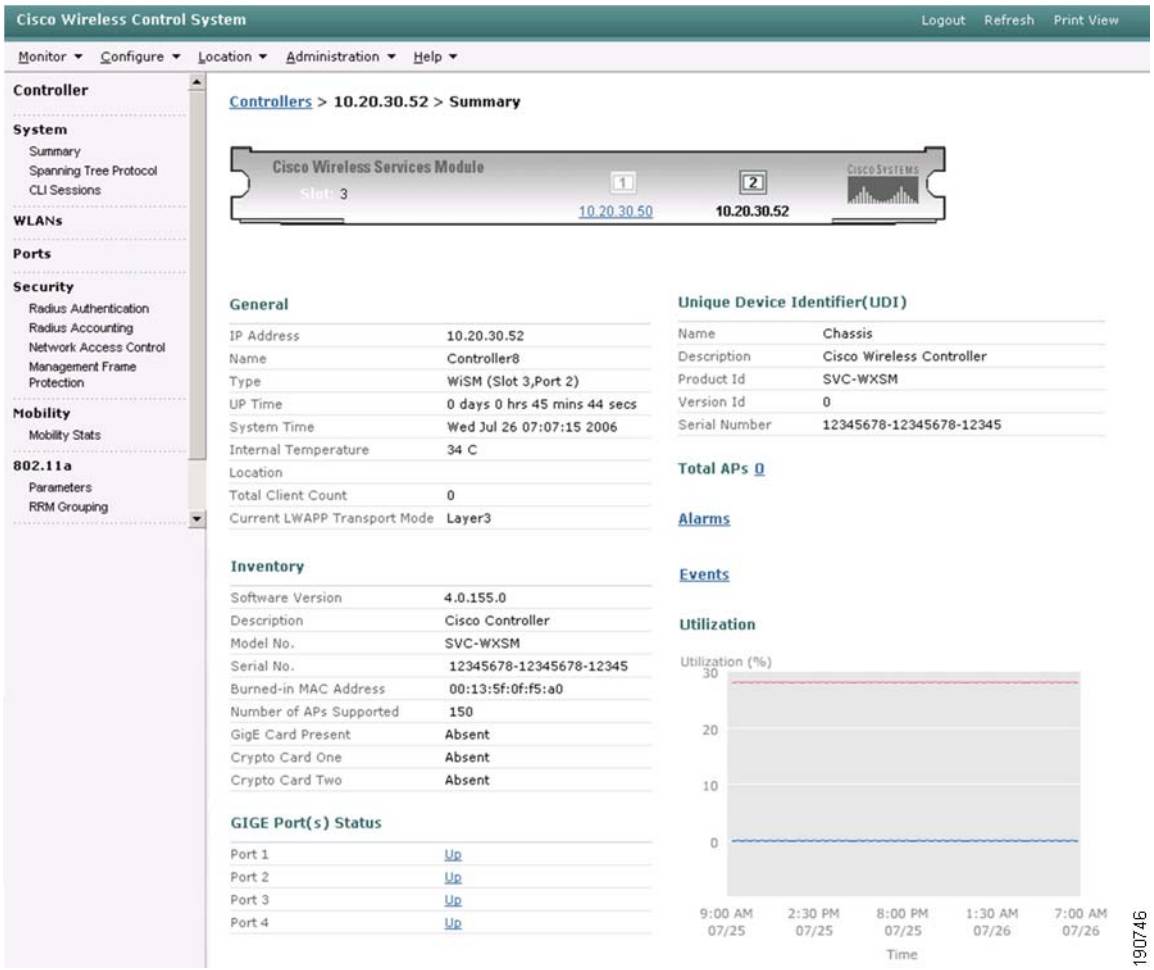

Note that in all cases, however, the screen format used in **Controller > Summary** is very similar with pertinent summary information in the main body of the screen and visual color representation used for the status of the physical Ethernet ports. Clicking on any of the red or green port graphics summons the **Monitor > Devices > Controllers > Ports** panel for the port in question, where a full range of Ethernet port performance information is available. Links are also available to applicable Alarms, Events, and AP Status pages from the controller summary display.

The left column of both Figure 8-34 and Figure 8-35 indicates the complete range of controller-specific information available from the controller summary page. WCS makes it possible to click on one category after another without requiring you to use the "back" browser function, thereby saving time and effort.

Monitoring Access Points

In addition to obtaining information about access point status indirectly via the **Monitor > Maps** and **Monitor > Devices > Controller** menu trees, WCS allows immediate and direct access the same information via the **Monitor > Devices > Access Points** menu selection. After clicking on **Monitor > Devices > Access Points**, you are presented with the **Access Point > Search Results** menu shown in Figure 8-36. Depending on the number of access points that WCS has discovered, the complete list of

search results menu may be quite long. Therefore, the "Search for APs" and "Select Radio Type" filters in the left-hand margin can be especially useful in narrowing down the total number of access points displayed.

Figure 8-36 Monitor > Devices > Access Points > Search Results

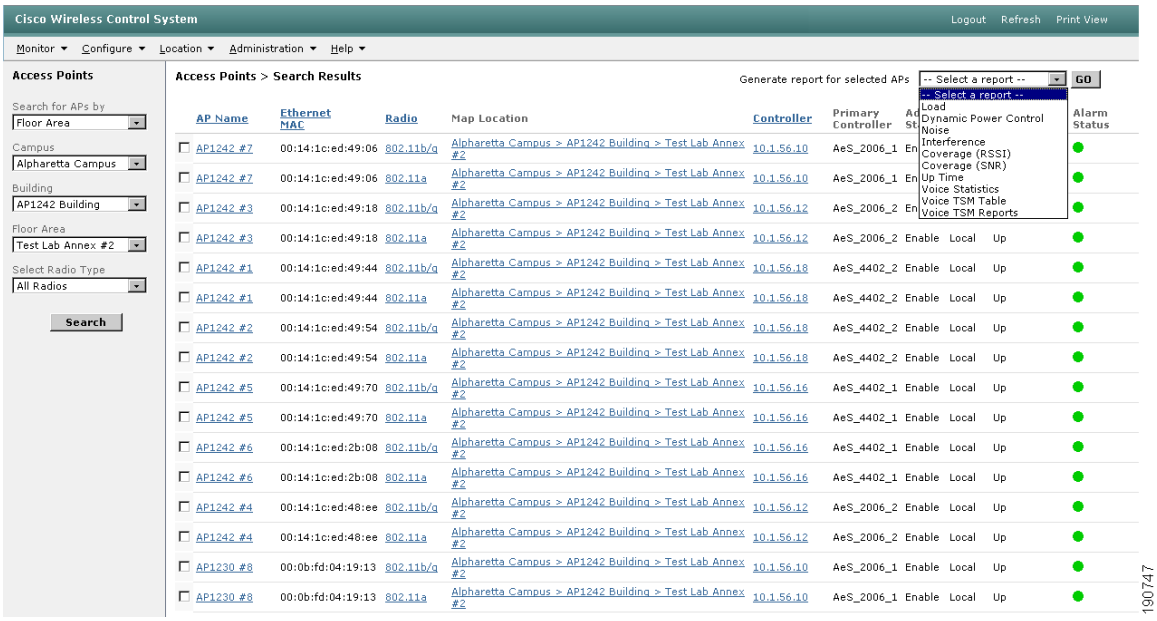

The various hyperlinks available from this screen allow access to a wide variety of information about the lightweight access point and its radio interfaces, the location map that it is currently assigned to, the controller to which it is currently registered as well as detailed information concerning the alarms that are currently active regarding it. Note that the Alarm Status indicator on this page is actually a hyperlink to detailed information about the alarm.

As was seen with **Configure > Access Points**, the search results page shown in Figure 8-36 displays *all* lightweight access points known to this WCS including any lightweight access points that are not currently registered to any WLAN controllers in the management domain. However, WCS *only* allows you to display detailed current status for lightweight access points that are currently registered to a WLAN controller (as indicated by the IP address of a WLAN controller appearing in the "Primary Controller" display column).

The drop-down command selector in the upper right-hand corner of Figure 8-36 indicates that several useful reporting functions are available from the search results page. Here you see the ability to generate load, dynamic power control, noise, interference, client distribution by RSSI / SNR, total uptime, and voice statistics, and traffic stream metrics reports for selected access points. This can be done by selecting up to five access point radio interfaces shown on the **Access Points > Search** Results screen, selecting a report type from the "Select a Report" menu drop-down at the top right-hand corner of the screen, and then clicking on **GO**.

Note All five selected access points *must* be displayed on the same screen to be selected for reporting (paging forward or backward for additional selections is not allowed). The uptime report allows only one access point to be selected.

Samples of the reports that are generated by these seven options can be seen in Appendix E, "Sample Monitor > Devices > Access Points Reports.".

Clicking on any of the AP Name hyperlinks yields the AP detail screen shown in Figure 8-37. This provides hyperlinks to the WLAN controller information found in **Monitor > Controllers** and the map information found in **Monitor > Maps** as well as hyperlinks to any alarms or events that concern this access point. On this screen, you can verify general operational parameters, software levels, access point model, serial number as well as the type of certificate with which it was provisioned. In addition, this page provides hyperlinks to 802.11a and 802.11b/g radio interfaces along with alarm status hyperlinks.

Figure 8-37 Access Point Detail

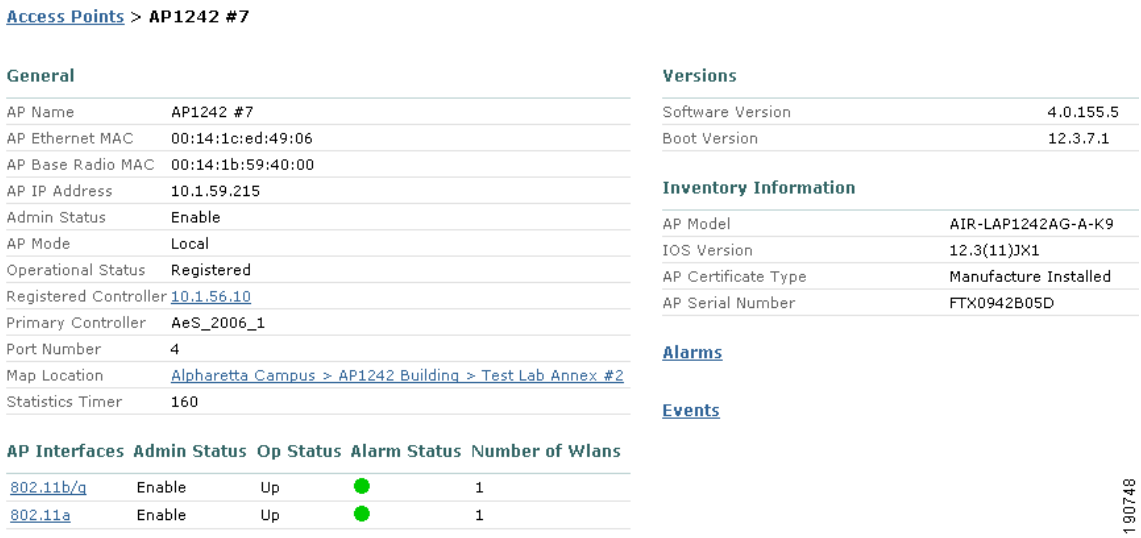

Clicking on one of the AP interface hyperlinks at the bottom of Figure 8-37 brings up the on-demand statistics page for the access point (Figure 8-38) where information about the status of the four radio resource management (RRM) profiles can be found (load, coverage, noise, and interference) along with information about the status of Management Frame Protection (MFP) on that access point. In addition, graphical charts displaying noise and interference by channel, percentage of client count, and channel radar detection along with receive, transmit, and channel utilization metrics can be found here and are illustrated in Figure 8-38. A bar-chart displaying the RSSI of neighboring lightweight access points as last detected by the lightweight access point you are monitoring is also shown.

Using the selector in the upper right-hand corner of the on-demand statistics page shown in Figure 8-38, considerable detail is available about access point operational parameters, 802.11 MAC counters, and any outstanding alarms and events associated with this access point.

Figure 8-38 Access Point On-Demand Statistics Display

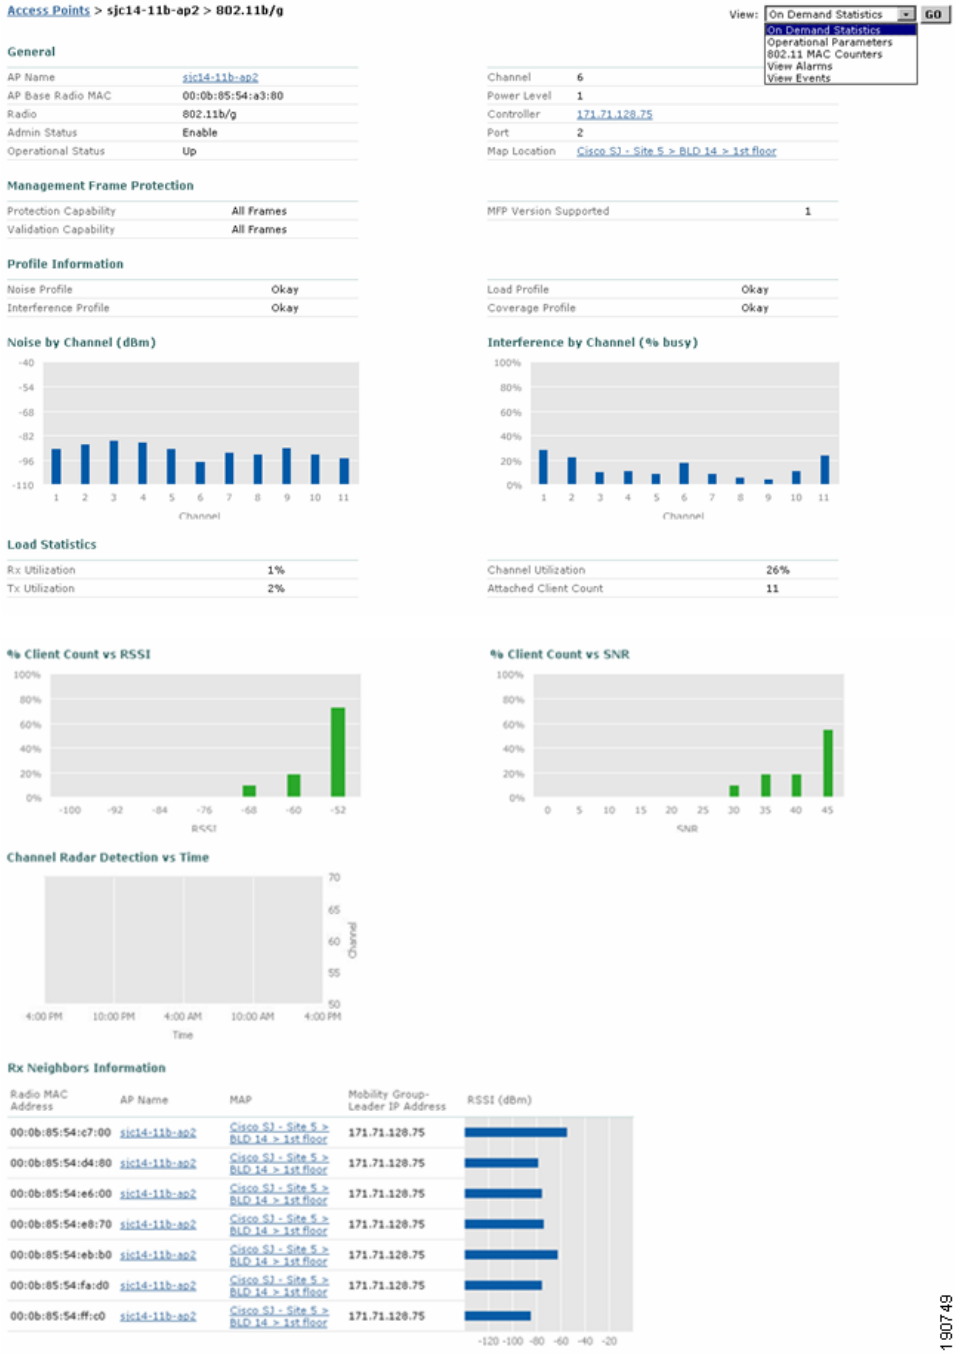

Monitoring Clients

As seen in the previous sections, WCS indirectly makes information available about clients associated to lightweight access points and controllers via the **Monitor > Devices > Controllers** and **Monitor > Devices > Access Points** menu options. WCS provides for direct access to an even larger base of client information via **Monitor > Devices > Clients**.

Although WCS is not a "client manager," it does allow you to configure and manage the various aspects of how client devices are permitted to interact with the network infrastructure. Through its tight integration with the various controllers, lightweight access points, and location servers comprising the wireless network, WCS accumulates a lot of information about the activities of clients on the network, regardless of whether they are probing, associated, or currently disassociated. This section briefly describes some of the information available via the **Monitor > Devices > Clients** selection.

After clicking on **Monitor > Devices > Clients,** the Clients Summary screen (seen in Figure 8-39) is immediately presented.

Figure 8-39 Monitor > Devices > Clients Summary

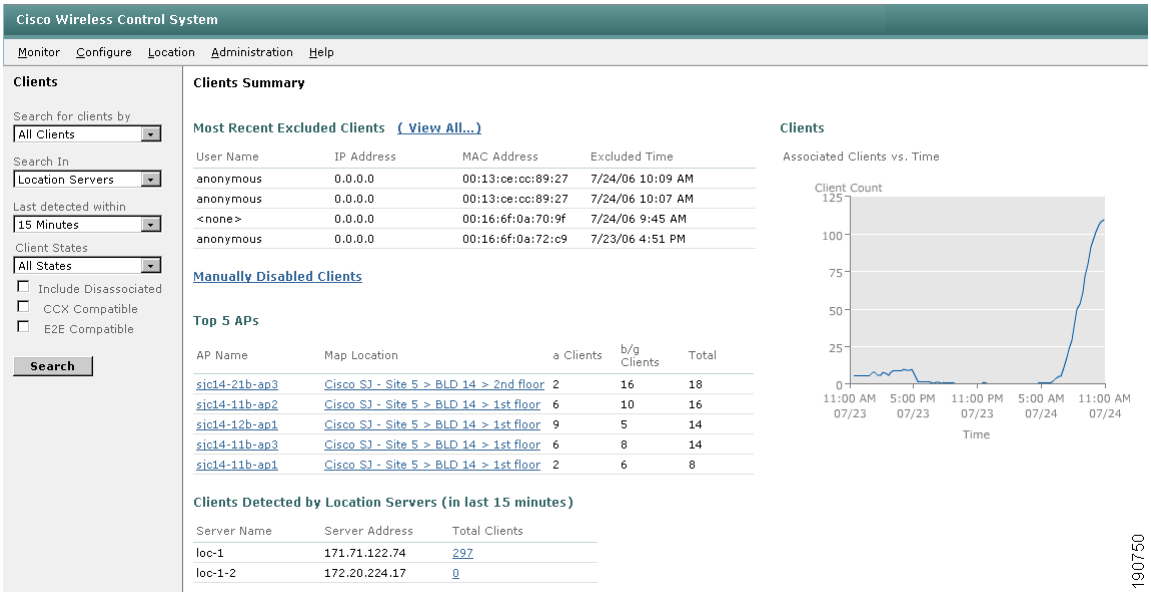

As seen in other summary screens, Client Summary attempts to display an overall view of client activity in your wireless LAN by displaying information in the following basic areas:

- Most Recent Excluded Clients—This area lists the clients that have been excluded from using the wireless LAN (also known as "blacklisted") because of the detection of one or more policy violations. For example, the client may have attempted association or 802.1x authentication and failed multiple times in succession. Security policies in the system normally exclude such a client from associating with the WLAN for a period of time as a security precaution against a possible intruder mounting an attack. Exclusion rules such as this are configurable via the WLAN template available at **Configure > Templates > WLAN**.

  Additional information regarding configuring client exclusion and the timeouts that are available can be found in the WCS menu bar under **Help > Online Help**.

  To view additional information about any of the excluded clients listed in this section, click on the **View All** hyperlink. This displays an Alarms screen where all the excluded clients are listed. The text of the reason for exclusion can be viewed by performing a mouse-over of the failure object name. Clicking on the failure object name hyperlink of any client displays detailed information about the alarm including hyperlinks to view the event history.

- Manually Disabled Clients—This hyperlink takes you to the **Configuration > Templates > Security > Disabled Clients** page where the current listing of clients that have been administratively excluded can be viewed.

- Top Five Access Points—This area refers to the same information as was shown on the Network Summary screen. See Network Summary, page 8-44 for further details.
- Clients Detected by Location Servers (with Fifteen Minutes)—This area indicates each location server that has been defined to WCS along with a hyperlink enumerating the total number of clients that have been detected within the last fifteen minutes. Note that if a location server is not defined to WCS, this area is still present; however, no entries for location servers are shown. Similarly, if location servers are defined but unreachable, the total client column shows a red "unreachable" alarm indicator.

  The total count of clients is a hyperlink that links the WCS user to the same information available by manually configuring the left-hand margin "Search For" feature. In this case, the hyperlink displays all clients found in all states in that particular location server within the last fifteen minutes.
- Associated Clients vs. Time—A graphical depiction of the number of associated clients plotted against time. By performing a mouse-over of various points on the graph, you can read the number of clients that were detected as associated at that particular sampling interval. The data points on the graph do not provide any hyperlink capability.

Control over the data displayed in **Monitor > Devices > Clients** is provided by the "Search For" selection bar in the left-hand margin. Figure 8-40 illustrates the full range of search criteria that are available to the WCS administrator when using this selection bar.

Figure 8-40 Search Criteria Available under Monitor > Devices >Clients

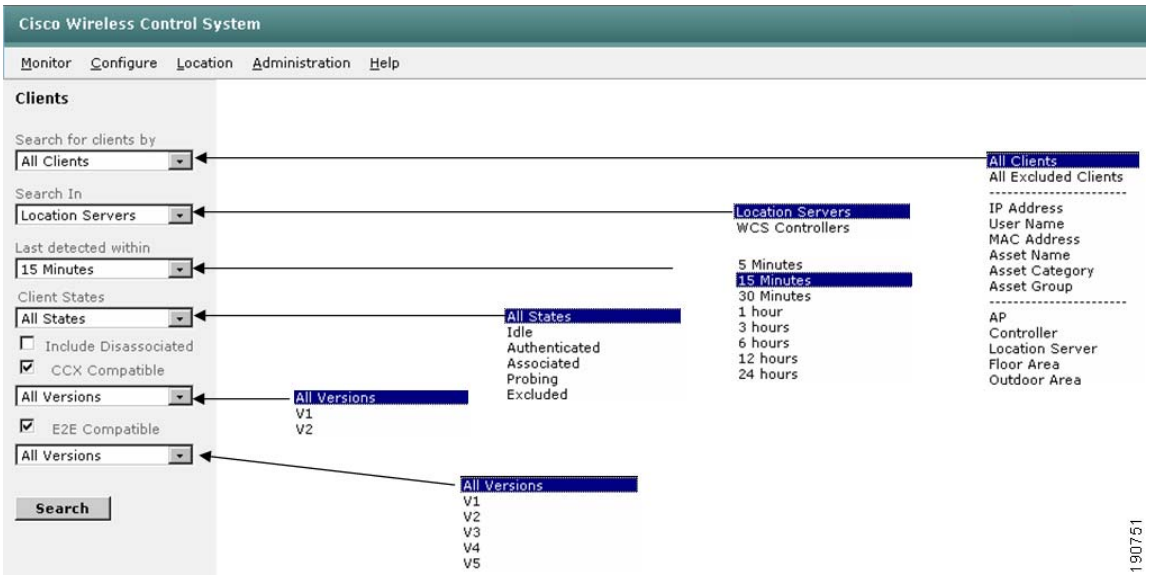

After selecting the appropriate criteria as shown in Figure 8-40, WCS returns a detailed listing of results in which several fields (user, AP, Map Location, and Link Test) are hyperlink-enabled as shown in Figure 8-41. Note the display of the miniature location floor map when performing a mouse-over of the user name hyperlink. Although small in size, this location floor map does indeed indicate the approximate location of the client. This handy method of quickly scanning each client in the list for their approximate location is available only when using a location-enabled version of WCS with the Wireless Location Appliance.

Figure 8-41 Monitor > Devices > Clients Listing

User	Vendor	IP Addr	MAC Addr	AP	Loc Server	802.11 State	SSID	Authenticated	Protocol	Map Location	
<none>	Unknown	0.0.0.0	00:02:8a:a2:2e:a0	00:0b:85:54:dc:b0	TME Loc2	Probing		No	802.11b	Cisco SJ - Site 5 Group>BLD 14>2nd floor	Link Test
one>	Unknown	0.0.0.0	00:02:8a:dc:40:74	00:0b:85:54:e5:70	TME Loc2	Probing		No	802.11b	Cisco SJ - Site 5 Group>BLD 14>3rd floor	Link Test
<no			e:44:1d:36	00:15:c7:a9:43:10	TME Loc2	Probing		No	802.11b	Cisco SJ - Site 5 Group>BLD 14>4th floor	Link Test
<no			e:45:65:e3	00:15:c7:a9:0b:20	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>4th floor	Link Test
<none>	Netgear	0.0.0.0	00:09:5b:2f:dc:9f	00:15:c7:a9:43:10	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>4th floor	Link Test
<none>	Netgear	0.0.0.0	00:09:5b:68:7f:c7	00:0b:85:54:e8:70	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>2nd floor	Link Test
<none>	Netgear	0.0.0.0	00:09:5b:a2:75:eb	00:0b:85:54:dc:b0	TME Loc2	Probing		No	802.11b	Cisco SJ - Site 5 Group>BLD 14>2nd floor	Link Test
<none>	Unknown	0.0.0.0	00:0a:5e:4b:74:a1	00:15:c7:a9:08:40	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>4th floor	Link Test
<none>	Unknown	0.0.0.0	00:0b:5f:6e:5b:74	00:0b:85:54:c7:00	TME Loc2	Probing		No	802.11b	Cisco SJ - Site 5 Group>BLD 14>1st floor	Link Test
<none>	Unknown	0.0.0.0	00:0b:fc:ff:af:30	00:0b:85:54:d1:c0	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>2nd floor	Link Test
<none>	Unknown	0.0.0.0	00:0b:fc:ff:af:50	00:0b:85:54:dc:b0	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>2nd floor	Link Test
<none>	Unknown	0.0.0.0	00:0b:fc:ff:af:b0	00:0b:85:54:ea:40	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>2nd floor	Link Test
<none>	Unknown	0.0.0.0	00:0b:fd:01:06:08	00:0b:85:54:e6:00	TME Loc2	Probing		No	802.11a	Cisco SJ - Site 5 Group>BLD 14>1st floor	Link Test
<none>	Intel	0.0.0.0	00:0c:f1:15:f3:94	00:15:c7:a9:42:60	TME Loc2	Probing		No	802.11b	Cisco SJ - Site 5 Group>BLD 14>4th floor	Link Test

190752

Clicking on the username hyperlink produces the Client Detail page (shown in Figure 8-42) that contains a wealth of information about the properties and statistics associated with the monitored client. Associated access point properties and security summary information for the client is also available from this page. If a location-enabled version of WCS is used with the location appliance present, a current floor map (which can be enlarged) showing the estimated location of the client is displayed under the Client Location heading.

Figure 8-42 Monitor > Devices > Clients Detail Page

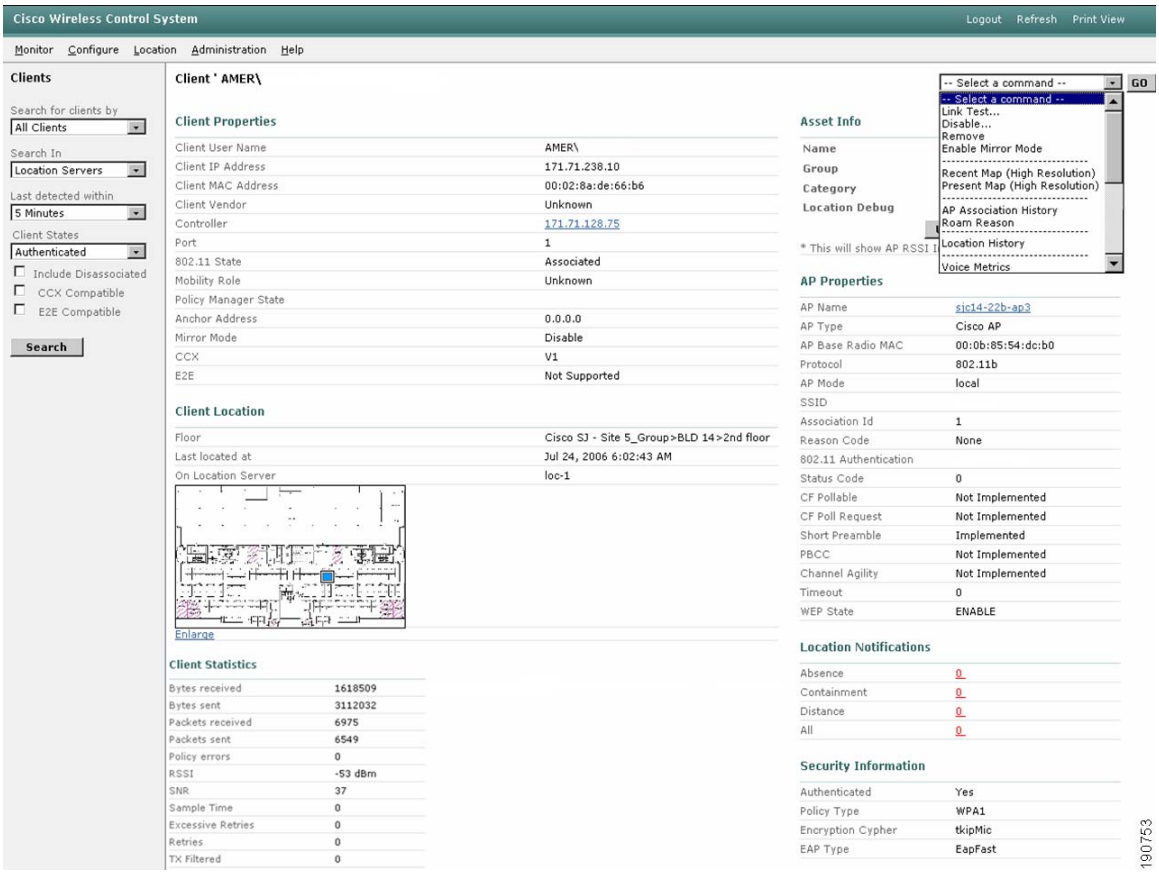

The location notification alarm display, AP name, and controller IP address are all hyperlinks leading to further detail about location notification alarms that may have been generated based on the movement of this client, the lightweight access point to which this client is associated, and the controller to which the lightweight access point is registered. The asset name, group, and category of the client device can also be updated in this location. Asset information entered for WLAN clients here can be used to define client filters in other WCS functions. For example, asset information can be used when large groups of WLAN clients must be quickly narrowed down by asset group or category and displayed as icons when viewing devices on location maps via the **Monitor > Maps** facility.

The client detail page also provides graphical trending displays for client RSSI, SNR, packets sent, and packets received, as shown in Figure 8-43.

Figure 8-43 Monitor > Devices > Clients Detail Graphical Displays

Using the command drop-down selector in the upper right-hand corner of the client detail screen in Figure 8-42, you can select from the following options:

- Link Test—Performs a test of the client wireless link and displays the results as shown in Figure 8-44.

Figure 8-44 Link Test Results

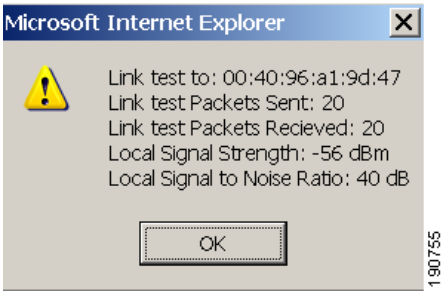

- Disable—Manually adds the client to the client exclusion list. Performing this action immediately de-authenticates/disassociates the client. Its newly-added presence on the exclusion list then prevents it from re-connecting.
- Remove—Disassociates/de-authenticates the client but does not place it on the exclusion list.
- Enable Mirror Mode—Enables this client as a candidate for the mirroring of all data originating from or destined to it. The data is mirrored to a spare Ethernet port on the WLAN controller that you select (you cannot use the service port interface for mirror mode). To use this function, you must enable port mirroring on the WLAN controller via **Configure > Controllers > Ports > General Config > Mirror Mode** and use an Ethernet protocol analyzer to capture the mirrored data. One method of accomplishing this is to connect the mirrored port on the controller to a standalone Ethernet switch (not part of the network to which the controller is already attached so as to prevent any spanning tree loops) and logically connect that switchport to a port on the standalone switch to which the Ethernet protocol analyzer is to attached, using the Cisco Catalyst Switched Port Analyzer (SPAN) feature. Information about how to configure the SPAN feature for various models of Catalyst switches can be found by performing a search at the following URL: http://www.cisco.com.
- Recent Map and Present Map—These two functions cause WCS (rather than the location appliance) to display the location of a client on the appropriate floormap, using either recent location history data or current client RSSI data. Note that when choosing to have WCS locate the client using current RSSI data, client wireless connectivity is briefly interrupted while the data is gathered. The client should reconnect and resume service with minimal disruption. For further information about using these two functions, see On-Demand Location of WLAN Clients, page 8-83.
- AP Association History—Selecting this function displays a graphical plot of the access points with which this client has associated versus time (shown in Figure 8-45). Note that there is a drop-down command selector in the upper right-hand corner that offers many of the same options just discussed as well as a link back to the client details screen.

Figure 8-45 Monitor > Clients > AP Association History

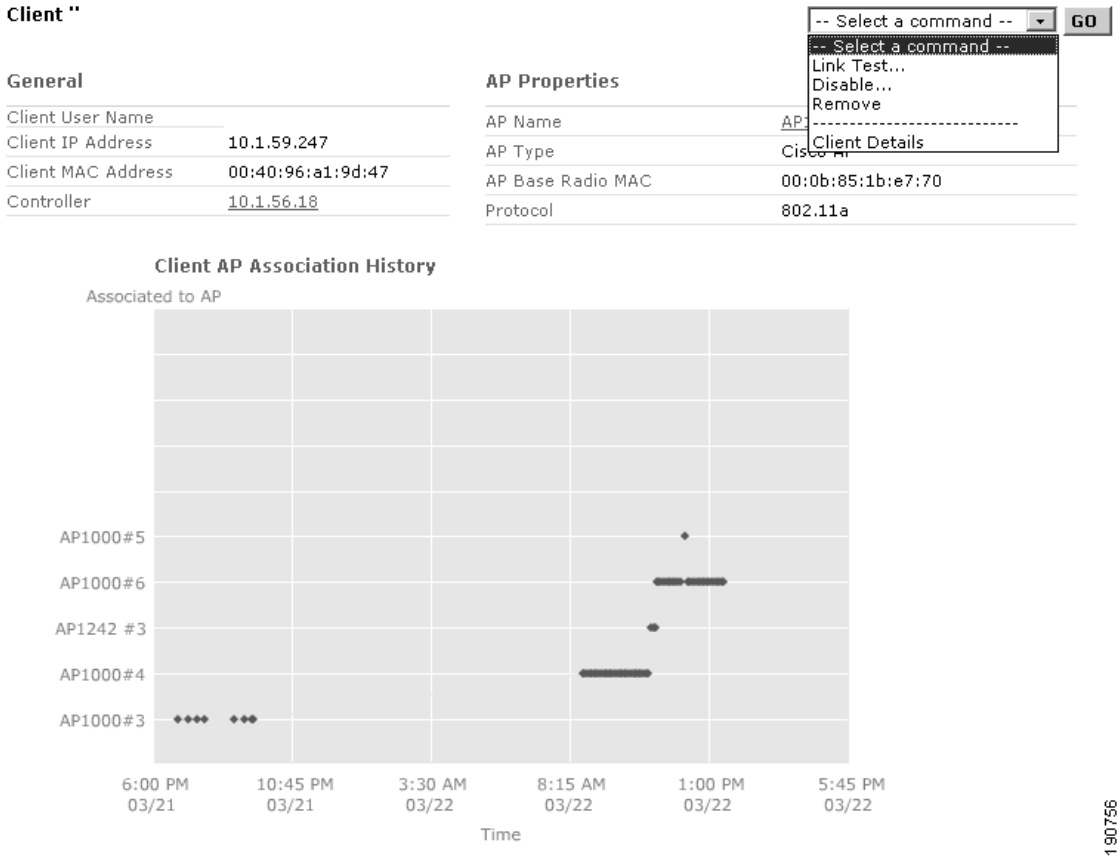

- Roam Reason—Using this option, information about when and why a client roamed within the environment can be found. A “Roam Reason” Report is generated that lists the following:
  - MAC addresses of the current as well as the immediately previous access point to which the client associated
  - The previous access point SSID and channel
  - The roam transition time
  - The reason why the client roamed
- Location History—This option appears with versions of WCS licensed for location use and is only functional when a location appliance is installed. It allows for the sequential display of the location history associated with a client device to better visualize and trace the movement of the client throughout the environment over time. This can be very useful, for example, in security and monitoring applications. WCS and the location appliance make it possible to view each location history record sequentially in this fashion, played back with a configurable time delay. The granularity of the “movement” shown depends on the interval with which client history records are recorded in the database. To see location history played back in this fashion, click the **Play** button as shown in Figure 8-46. Past location history should be displayed both in tabular and graphical form. Large amounts of location history data may be more readily viewed by reducing the “Change Selection Every” interval from 2 seconds to 1 second.

  For additional details about location history features and the location appliance, see *Wi-Fi Location Based Services—Design and Deployment Considerations* available at http://www.cisco.com.

Figure 8-46 Monitor > Clients > Client Detail > Location History

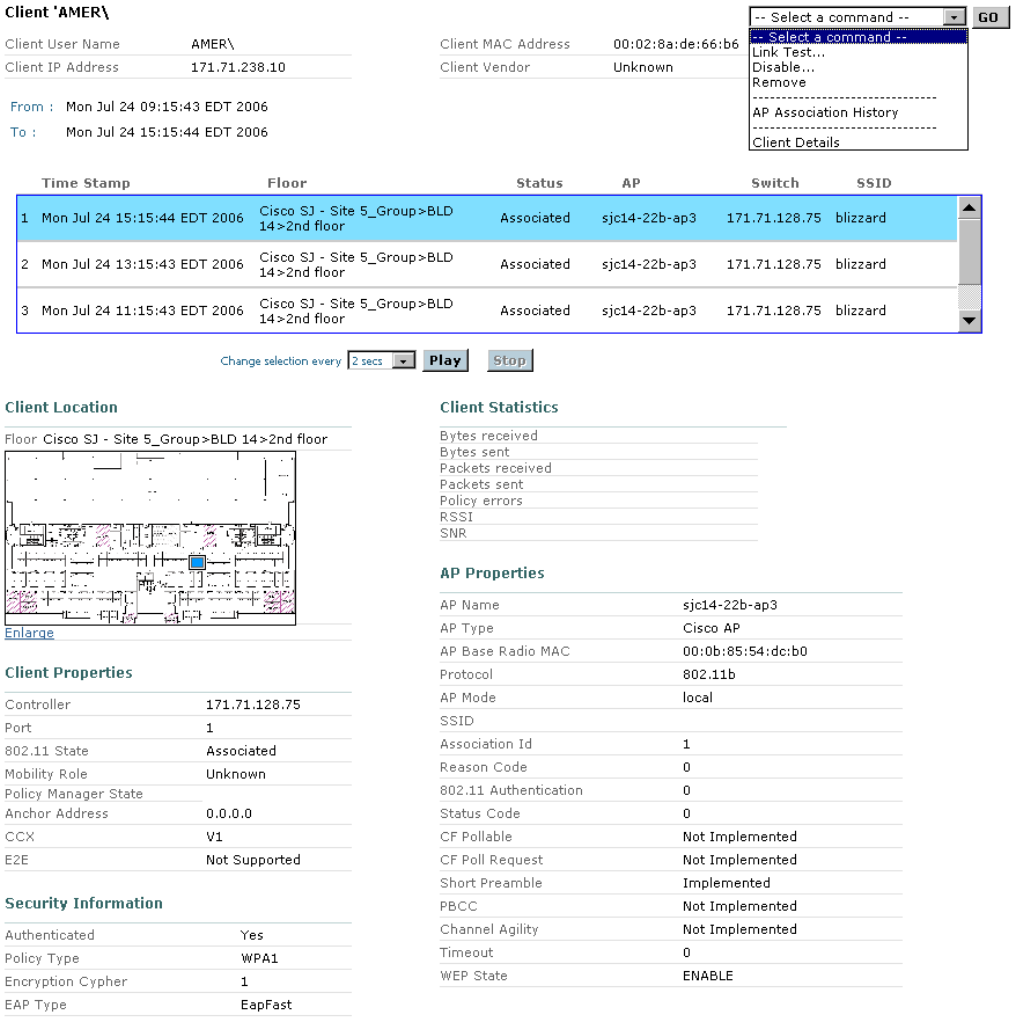

- Voice Metrics—This displays the voice stream metrics report and requires that voice traffic stream metrics (**Configure > Controller > ipaddress > 802.11bg > Voice Parameters > Enable Voice Metrics**) be enabled on the WLAN controller. See the *Cisco Wireless LAN Controller Configuration Guide, Release 4.0* and *Cisco Wireless Control System Configuration Guide, Release 4.0* for further information about enabling voice stream metrics and voice metric reporting.

Monitoring Asset Tags

Information about asset tags that have been detected by lightweight access points and controllers can be accessed indirectly via the **Monitor > Controllers** and **Monitor > Access Points** menu selections. When WCS has been installed and licensed for location-based services, the WCS user is able to directly access this information via the **Monitor > Devices > Tags** function. Note that the **Monitor > Devices > Tags** submenu does not appear for a WCS that has been licensed only for base level functionality (WCS-Base). Only 802.11 active RFID Layer 2 asset tags (such as those from AeroScout) are displayed under **Monitor > Devices > Tags**. Asset tags acting in Layer 2 mode typically do not associate or authenticate to the WLAN but rather communicate their payload information via Layer 2 multicasts. Other types of

802.11 active RFID asset tags that associate/authenticate to the wireless infrastructure are detected as WLAN clients and are not listed under **Monitor > Devices > Tags** (instead, they can be found under **Monitor > Devices > Clients**).

Note For a detailed discussion of Layer 2 and other types of asset tag technologies, see *Wi-Fi Location Based Services—Design and Deployment Considerations* at the following URL: http://www.cisco.com.

After clicking on **Monitor > Devices > Tags,** the Tag Summary screen (seen in Figure 8-47) displays the total number of asset tags detected by the location appliance within the previous fifteen minutes as a hyperlink. Clicking on this hyperlink allows you to quickly initiate a search in the associated location server for all tags from all vendors detected within the last fifteen minutes.

Figure 8-47 Monitor > Tags Tag Summary Screen

If you want to modify the terms of the search (that is, to specify a different search time frame, vendor, or other tag search criteria) use the drop-down menus in the left-hand margin of Figure 8-47.

The tag search results page, tag detail page, and asset tag location history display all have a very similar look and feel to their respective client page counterparts discussed in the previous section. On the Tag Search Results page, you can see tag-specific information such as tag asset information, tag vendor, location server, and tag battery status. Of the displayed fields, tag MAC address, switch IP address, and tag map location are hyperlink-enabled.

Note that performing a mouse-over of the tag MAC addresses shows a miniature location map with the asset tag located estimated using a yellow tag icon, as shown in Figure 8-48.

Figure 8-48 Tag Search Results

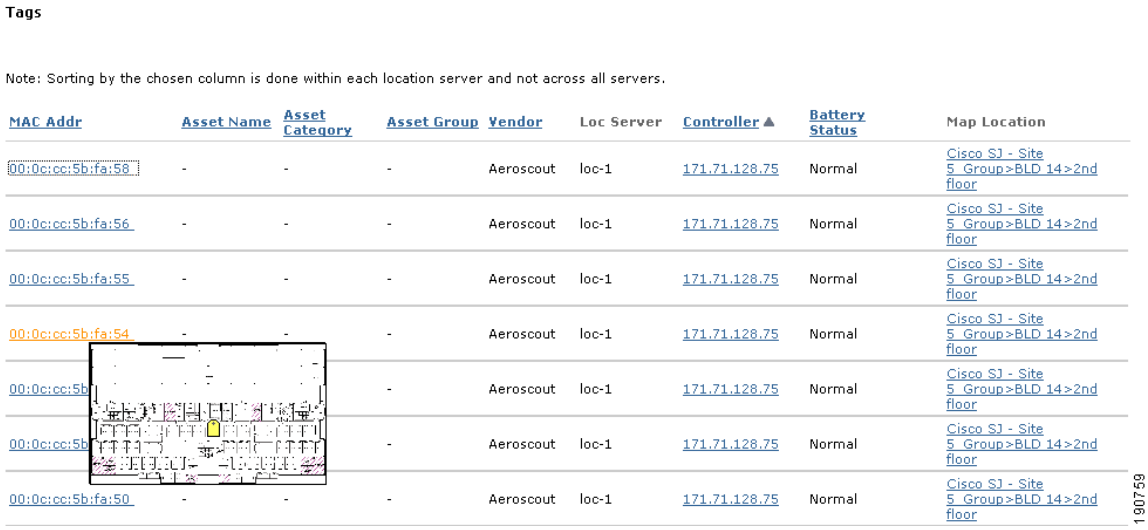

The tag detail page is similar to what has been seen before for clients, and allows for only the Location History option via the right-hand command drop-down menu selector (see Figure 8-49).

Figure 8-49 Tag Detail Page

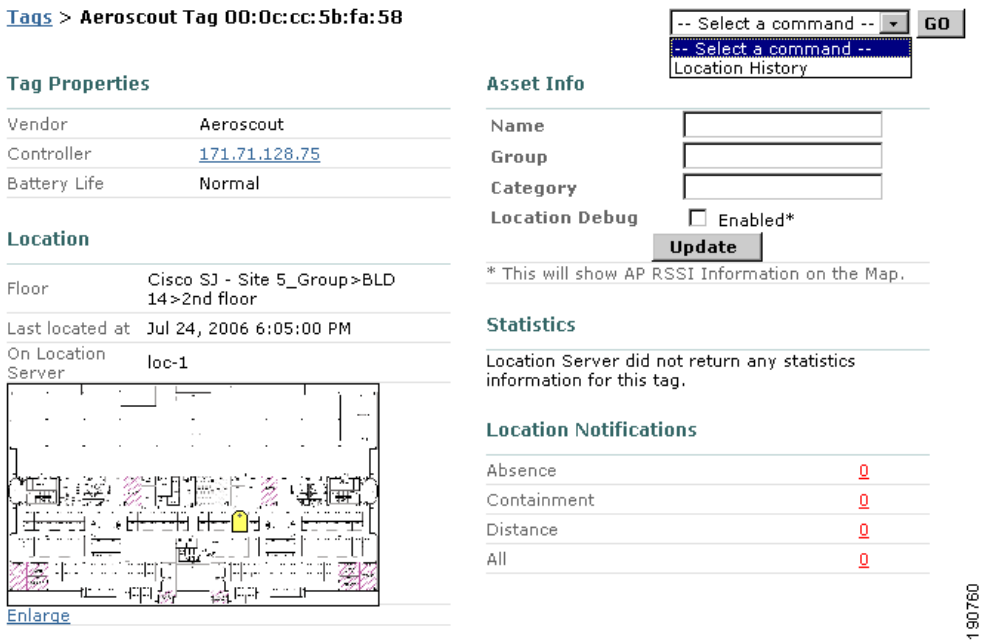

The format and function of the asset tag location history display is similar to that described for clients as well. Note the check box for "Location Debug" that enables WCS to display graphically the RSSI at which each access point has detected the asset along with an indication of how much time has passed since the asset tag has been detected. The enabling of Location Debug and how it can be used to facilitate the proper design and maintenance of a location-aware Cisco WLAN is discussed further in *Wi-Fi Location Based Services—Design and Deployment Considerations*.

Monitoring Security

WCS provides a summarization of security-related events in the network via the Security Summary page (shown in Figure 8-50), which is available from the main menu bar via **Monitor > Security**. Although every page in WCS displays the latest tally of critical, major, and minor alarms via the lower left-hand corner (the alarm monitor), the Security Summary page provides an especially detailed view of the security-related events that have recently transpired in the wireless network.

Figure 8-50 Monitor > Security Summary Page with Main Information Groups Emphasized

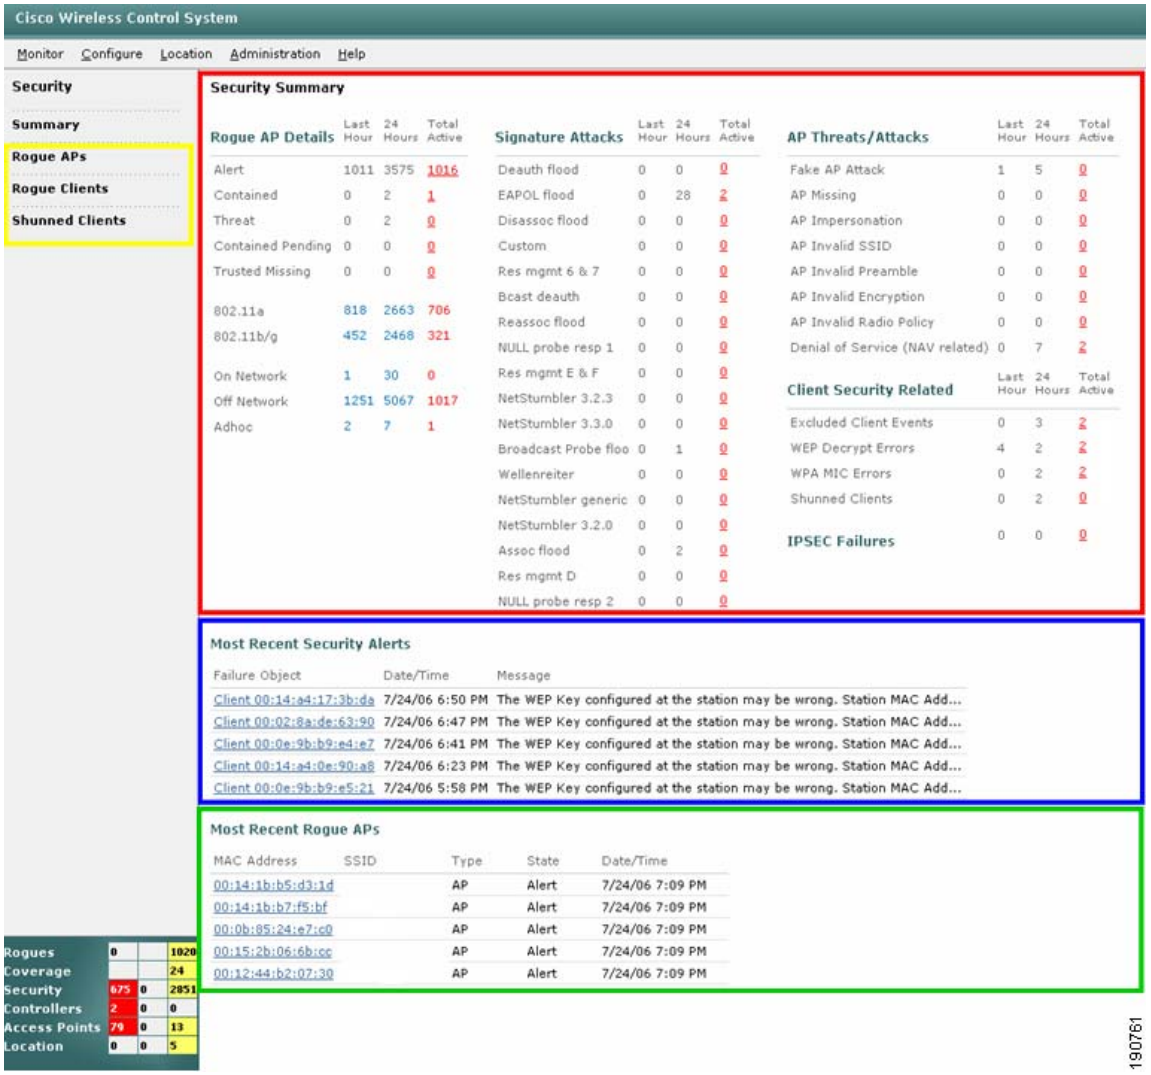

Because consumer-grade 802.11 access points are so readily available, maintaining constant vigilance against the proliferation of unauthorized rogue access points and rogue clients in the enterprise is typically one of the top priorities of security staff. Employees, contractors, and sometimes even visiting customers and guests commonly plug these unauthorized access points into existing LANs or build *ad hoc* wireless networks to facilitate their own mobility without the knowledge or consent of corporate IT or security departments.

These rogue access points can be a serious breach of network security because they can be very easily plugged into a network port behind the corporate firewall. Because these products often ship configured for easy wireless connectivity, security is usually disabled, and their owners often leave these settings at factory defaults. This being the case, it is very easy for malicious third-party users or outside hackers to gain access to the corporate intranet by using these unauthorized rogue access points as an easy entry point to the corporate intranet. After being discovered, the location of these unsecured portals can be published on the Internet, thereby drawing even greater attention to them from an unscrupulous community and increasing the odds of an enterprise security breach.

Rather than having a technician with a wireless analyzer constantly patrolling physical sites for new rogue access points, the Cisco Unified Wireless Network offers the ability to automatically collect information on these unauthorized devices via its managed access points. This allows the system administrator to determination the location of these rogues and make conscious decisions about their status.

These integrated anti-rogue client and access point capabilities allow WCS system administrators and other authorized users to do the following:

- Receive new rogue access point and client notifications and establish the location of these rogues, eliminating the need for a technician to periodically visit each site with a wireless network analyzer
- Monitor unknown rogue access points until they are eliminated or consciously acknowledged as benign
- Initiate containment actions against rogue access points and their clients by sending the clients deauthenticate and disassociate messages, discouraging further communication
- Acknowledge the presence of benign rogue access points when they are not attached to the enterprise wired LAN and can be seen to be located outside of the physical premises of the enterprise
- Acknowledge the presence of benign rogue access points when they are not attached to the enterprise wired LAN and can be seen to be located within the physical premises of the enterprise. These types of rogues typically operate in testing labs or other standalone networks that are not attached to the enterprise intranet. Rogue operation of this nature is typically for legitimate business purposes and is performed under the full knowledge and approval of corporate security departments.
- Tag other rogue access points as unknown until they are eliminated or acknowledged
- Tag rogue access points as contained, which implies that the containment process has been successful

Figure 8-50 shows the Security Summary page divided into four key areas:

- *Fast access hyperlinks* (shown within the yellow rectangular area in the left-hand column):
  - Listing of rogue access point alarms of all severities and categories

    See Monitoring Events, page 8-75 for a further discussion of the rogue access point alarm detail page.
  - Rogue Client search screen

    To use this facility, the search criteria in the left-hand column of the Rogue Client search screen must be configured to search for desired rogue clients in either the databases of WCS or the location appliance. When this is performed, a listing of rogue clients such as that displayed in Figure 8-52 can be viewed.

Figure 8-51 Monitor > Security > Rogue Clients

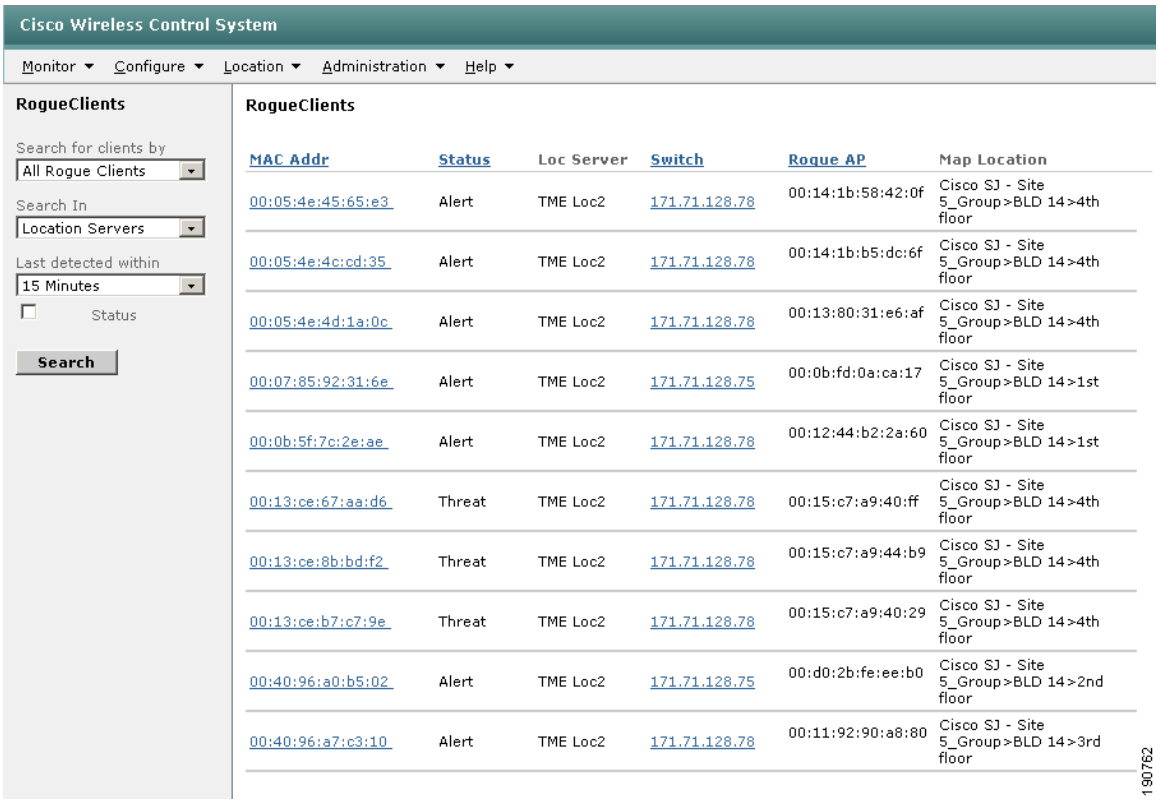

Clicking on any of the rogue client MAC addresses yields the rogue client detail screen (shown in Figure 8-52) providing detailed information about the rogue client such as which and how many lightweight infrastructure access points have detected it, when it was first and last heard, and its location.

Figure 8-52 Rogue Client Detail

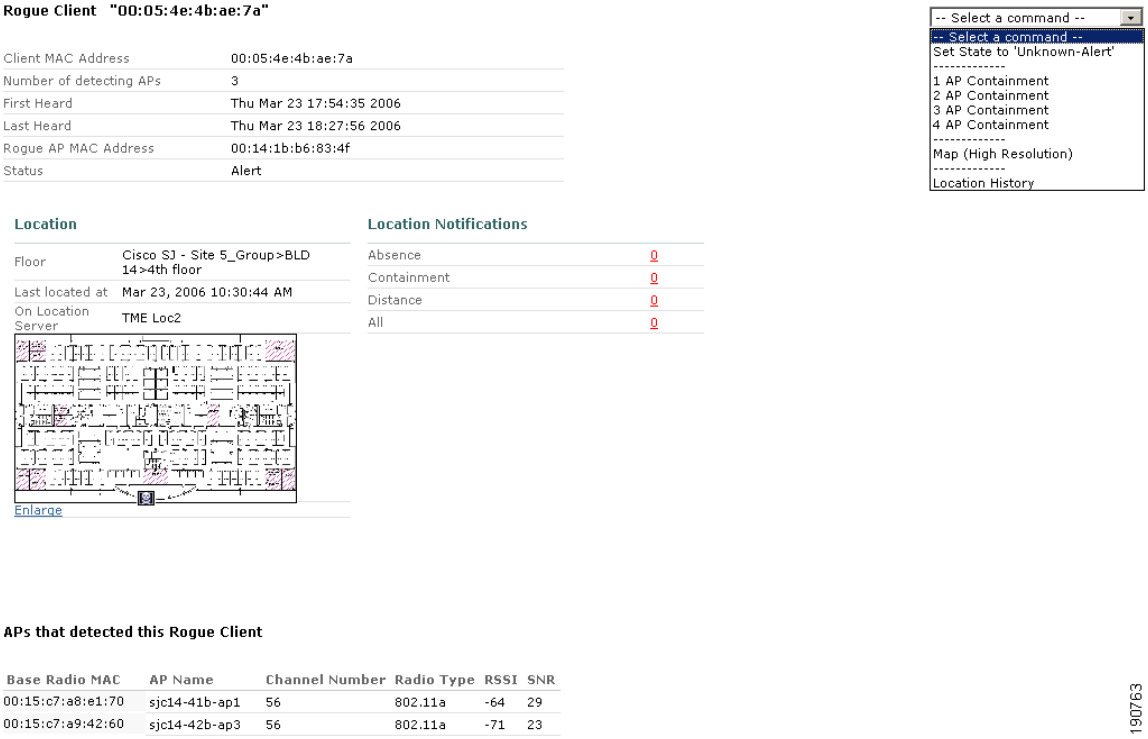

The drop-down menu selector in the upper right-hand portion of the rogue client detail screen shown in Figure 8-52 allows the WCS user to perform additional functions concerning the rogue client such as displaying its location history and current location, changing its status, or initiating containment of the rogue client. When you select level 1 containment, one lightweight infrastructure access point in the vicinity of the rogue client sends de-authenticate and disassociate frames to the client. When you select level 2 through 4 containment, two through four lightweight infrastructure access points participate in containing the rogue client.

- The Shunned Clients search screen.

  Using this facility, a search can be conducted of all associated clients that have been detected by a Cisco IDS device as sending malicious traffic through the network. The IDS device detects such activity and sends “shun” requests to Cisco Wireless LAN Controllers, which in turn disassociate the client device. This search facility provides access to a listing of shunned clients by client IP address, IDS sensor address, and controller.

- *Security Alarms*—This area typically occupies the upper half of the security summary page (shown in Figure 8-50 within the red rectangular area). It is subdivided into subgroups that provide last hour, 24 hour, and active alarm counts pertaining to rogue access points, signature attacks, AP threats/attacks, IPsec failures, and client-related security alarms. The total active alarm counters in this area are hyperlinks to a common list of alarms that pertain to the particular category. For example, clicking on the total active alarms counter for rogue AP alerts takes you to the appropriate listing of alarms.

Note Note that there is no implied relationship of (Last hour alarms) + (Last 24 hour alarms) = Total Active Alarms. Last hour and last 24 hour alarm counters are a tally of *all* alarms received in those time frames. The Total Active Alarms counter is the sum total of alarms received *net any corresponding clear alarms received*. Therefore, it is not unusual to see (Last hour alarms) + (Last 24 hour alarms) > Total Active Alarms.

Most of the categories in the Security Alarms area are self-explanatory, but a few are clarified as follows:

- *Threats*—Any rogue APs detected as being on the same wired network as your infrastructure lightweight access points are considered a threat. The detection of a threat is always considered a critical event/alarm.
- *Contained Pending*—This is a transitory state for a rogue AP that is in the process of being contained.
- *Trusted Missing*—A rogue AP that has been marked as Known Internal or Known External but is now determined to be missing from the management domain.

- *Most Recent Security Alerts*—This area is shown within the blue rectangle in Figure 8-50 and displays up to five of the most recent security alerts detected. The entire text of each alert can be seen by simply performing a mouse-over of the listed failure objects. Each failure object is itself a hyperlink to the associated alarm detail page.
- *Most Recent Rogue APs*—Shown within the green rectangular area in Figure 8-50, this portion of the Security Summary page displays up to five of the most recent rogue APs detected. Details are displayed during a mouse-over of each MAC address, and each address is a hyperlink to the Alarm Details page associated with the rogue AP.

Monitoring Events and Alarms, and Generating Notifications

WCS contains an alarm and notification subsystem that maintains an automated watch over the network, alerting the administrator (or anyone else with an e-mail account) even when they are not logged into WCS. Multiple recipients can be alerted about potential issues as they are discovered and before they turn into full-fledged problems.

Relationship Between Traps, Events, Alarms, and Notifications

In the Cisco Unified Wireless Network:

- The WLAN controllers and lightweight access points monitor activity that involves RF.
- Controllers monitor the status of lightweight access points and enforce defined policies.
- Location appliances monitor the movement of clients and assets.
- WCS monitors all the location servers, controllers, and lightweight access points that are within its management domain.

If operational exceptions (access point not found, controller unreachable, location appliance not responding, and so on) or other changes in state are deemed to have taken place within the domain of control of the controller (excessive interference, rogue access point detected, and so on), an *event* is registered as having occurred within the WLAN controller. An event is an occurrence or detection of some condition in and around the network. It can be a report about radio interference crossing a threshold, the detection of a new rogue access point, a controller rebooting, or some situation which takes places at a particular time.

Events are communicated to WCS via one of two ways:

- SNMP traps
- Normally scheduled polling

SNMP traps enable WLAN controllers and location appliances to notify WCS of significant changes in condition within the network by way of an unsolicited SNMP message known as a *trap*. Traps are normally sent to *trap receivers* at UDP port 162. Controllers can be configured to send traps to multiple trap receivers. The trap port to which a controller sends traps cannot be changed via WCS (it must be changed via the controller web interface or CLI). Traps are sent to the same trap port for all trap receivers defined in a controller.

Controller traps can be enabled or disabled in WCS via **Configure > Controllers > Management > Trap Control**.

Depending on the nature of the condition detected by the controller, a system log (*syslog*) message may also be generated. An example of this is the automatic addition of a client to an exclusion list because of repeated authentication failures or the loss of connectivity to a lightweight access point that was formerly associated with this controller. Note that traps and syslog messages do not necessarily duplicate one another. Events may not always be severe enough to generate syslog messages but may trigger pre-configured traps, and conditions that are severe enough to trigger a message to be sent to the system log may not be addressed by any configured traps.

SNMP traps and UDP syslog messages are logged internally at the controller and can be sent to external destinations. Both SNMP traps and syslog messages are sent using UDP; controllers can send traps to multiple IP destinations but they send syslog messages only to a single IP destination. Traps should always be sent to WCS for proper management function, but there is no need to send syslog messages to WCS because WCS does not natively run a remote syslog server to process them.

Figure 8-53 illustrates the various individual trap categories for a WLAN controller.

Figure 8-53 Configure > Controllers > Management > Trap Control

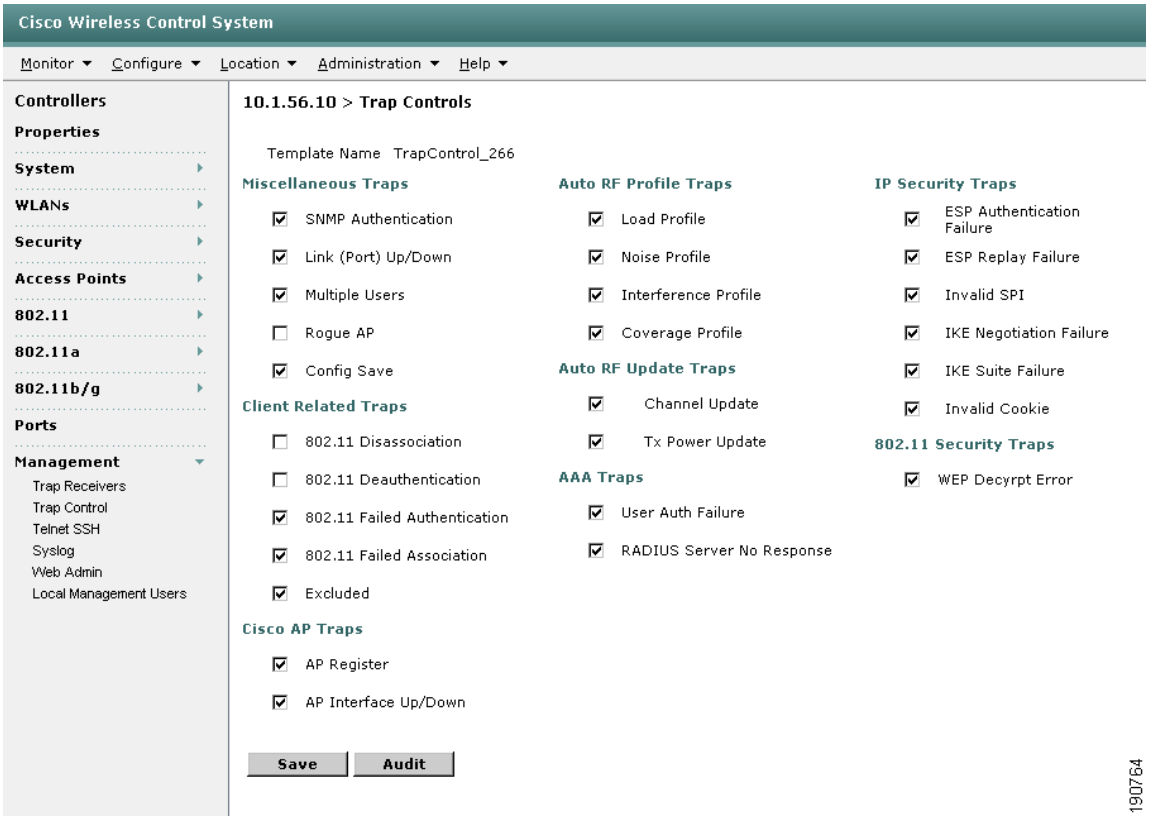

Trap categories are as follows:

- Miscellaneous traps
  - Multiple Users Traps—When enabled, this trap notifies all trap receivers when more one user logs into the CLI of a WLAN controller using a particular user credential. The value for this threshold is hardcoded in the controller and is currently set to one. When WCS receives this trap from the WLAN controller, a critical event is logged as shown in Appendix B, "WCS Event and Alarm Severities.".
- Auto RF profile traps
  - Load, noise, interference, and coverage profiles—When enabled, these traps notify all trap receivers when 802.11a or 802.11b/g load, noise, interference, and coverage threshold violations have been detected by any RF interface of any lightweight access point registered with the WLAN controller. These profiles are specified in WCS using **Configure > Controllers > 802.11a > RRM Thresholds** and **Configure > Controllers > 802.11b/g > RRM Thresholds.** Thresholds can also be specified using Policy Templates as discussed in Defining and Applying Policy Templates, page 8-22. When WCS is notified that one or more of these thresholds have been violated, a minor severity event is logged as shown in Appendix B, "WCS Event and Alarm Severities.".

In addition to the traps defined in Figure 8-53, there are a few others that are associated with thresholds to consider. The thresholds associated with these traps are not configurable:

- Too Many Unsuccessful Login Attempts—This trap notifies all trap receivers when a user using the WLAN controller CLI fails to successfully login after five attempts. When WCS receives this trap from the WLAN controller, a critical event is logged as shown in Appendix B, "WCS Event and Alarm Severities.".

- Maximum Rogue Count Exceeded—This trap notifies all trap receivers when the total number of rogue APs detected exceeded certain prescribed thresholds. The thresholds are as follows:
  - For WLCM and 2006 WLAN controllers—A maximum of 30 rogue APs detected per infrastructure AP; maximum of 125 rogue APs detected per WLAN controller.
  - For all other WLAN controllers including WiSM—A maximum of 30 rogue APs detected per infrastructure AP; maximum of 625 rogue APs detected per WLAN controller.

When WCS receives any of these traps from the WLAN controller, a critical event referring to a potential "Fake AP or other attack" is logged as shown in Appendix B, "WCS Event and Alarm Severities.".

WCS should always be configured as a trap receiver to ensure proper function. Trap receivers may be defined and removed for controllers via the **Configure > Controllers > Management > Trap Receivers** menu selections. WLAN controllers allow the definition of up to six trap receivers, and WCS allows the application of up to six trap receiver templates to a WLAN controller.

Appendix D, "Examples of SNMP Traps," contains several examples of SNMP traps captured during lab testing. Although not an exhaustive list of all traps available from Cisco WLAN controllers, this appendix does show actual received traps and decodes much of their content for easier viewing. Complete trap definitions for 4400, 4100, and 2000 Series WLAN controllers are defined in MIB files that are available to registered users on the Cisco Connection Online (CCO) at http://www.cisco.com.

Although WCS itself is not a syslog server, it does enable the configuration of syslog receivers for controllers via **Configure > Controllers > Management > Syslog**. If you have a system in your network that can function as a syslog server and accept remote syslog updates, you can define that system as the recipient of syslog messages from WLAN controllers using this feature.

Note Syslog message error level (that is, Critical, Error, Informational, and so on) is not configurable via WCS but can be set via the controller GUI.

All syslog messages are sent to UDP port 514 using facility level Local0; therefore, you should ensure that your syslog server is configured appropriately for UDP port 514. Only UDP is supported as a transport for syslog traffic at this time.

After receiving and processing a received trap, WCS logs the fact that an event has occurred at the trap originator. Events are filtered into assigned classes of severities. These events trigger *alarms* of corresponding severity at WCS. An alarm is a WCS response to one or more related events. If an event is considered of high enough severity (critical, major, minor warning, clear, or informational), the WCS raises an alarm until the condition that resulted in the alarm is judged to be no longer occurring. For example, an alarm may be raised while a rogue access point is detected, but the alarm terminates after the rogue has not been detected for several hours. A WCS administrator currently has no control over which events generate alarms, when they time out, or what severity they are.

One or more events can result in a single alarm being raised. Certain classes of alarms are deemed of utmost importance and when they occur, WCS can generate a outbound notification to alert critical personnel that events have occurred that may warrant their immediate attention.

Figure 8-54 illustrates the logic behind this event-driven alarm and notification system in more detail. This example is that of a WLAN controller, but this same process is used by all members of the Cisco Unified Wireless Network that wish to enlist the alarming and notification services of WCS.

Figure 8-54 High-Level WCS Event, Alarm, and Notification Flow Diagram

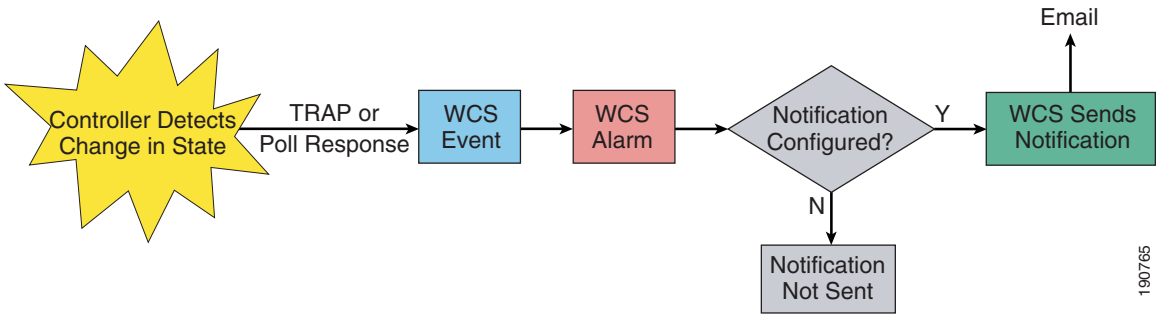

The following sequence takes place:

1. The controller detects that a change in state has occurred within the portion of the wireless network that the controller and its lightweight access points are monitoring. This can be the sudden unavailability of a network resource such as one of the lightweight infrastructure access points registered with the controller, or the identification of a potential security breach such as the detection of a rogue access point. Controllers can send information about such state changes in their environment to WCS, either via an unsolicited SNMP trap, a poll response, or both.
2. On reception of the trap, WCS registers an event and the event is filtered according to one of the following five severities: critical, major, minor, clear, or informational. Appendix C, "Example of Wireless LAN Controller Initial Setup," contains a listing of the various event and alarm messages that comprise these severity classfications.

Note There is actually a *sixth* severity class for "warning" events and alarms. However, as of this writing, the warning severity class is reserved for potential future use. There are no traps, events, or alarms currently classified as warnings.

- As Figure 8-54 illustrates, the reception of a trap or a poll response leads WCS to log an event that in turn can lead to the triggering of an alarm. It is very important to distinguish WCS *events* (which represent occurrences or a change in condition on a network device) from WCS *alarms* (which are states that occur only on WCS that are the direct result of an event). The alarm state can be cleared either manually or by an administrator. It can also be cleared automatically by an event indicating that the condition responsible for the original alarm state for a managed device has been resolved.

  Active alarms (that is, alarm status of other than "clear") have an indefinite lifetime while cleared alarms linger within WCS for 24 hours. Events, on the other hand, remain in the WCS database for up to seven days. The current set of active and cleared alarms can be viewed at **Monitor > Alarms** while the WCS event log can be viewed at **Monitor > Events** (these are examined in more detail in Monitoring Events, page 8-75 and Monitoring Alarms and Configuring E-mail Notifications, page 8-76).

  Alarms are assigned to severity classes in a similar fashion to events. For the majority of non-critical alarm severities, WCS logs the alarm, changes displayed icon colors appropriately, and increments the alarm monitor counters displayed in the lower left-hand corner of each WCS screen. Indication of the alarm can be seen on the **Monitor > Alarms** screen and by changes in the color of icons seen on campus, building, and floor maps; pictorial displays of network equipment; and other screens within WCS.

- Alarms classified as critical have the potential for immediate service impact and as such offer the option for external notification of key personnel. (As is discussed subsequently, coverage hole alarms are a special case in that they also provide the option for external notification even though they are not classified as “critical alarms”.)

  Recipients of WCS-generated notifications commonly include organizational team members such as network administrators, lead technicians, operations management, and perhaps those members of management responsible for mission-critical applications that depend on wireless services. You can see from the event flow in Figure 8-54 that if notifications have been enabled and properly configured, WCS ultimately issues an email notification. This is performed via one or more Simple Mail Transfer Protocol (SMTP) servers defined to WCS.

  In determining whether an e-mail notification should be sent when a particular alarm is triggered, WCS uses alarm severity in conjunction with alarm categories. There are currently the following seven alarm categories:

  – Rogue detection
  – Coverage holes
  – Security
  – Access points
  – Controllers (switches)

  Note The use of the term “switches” here is not intended to refer to Ethernet LAN switches such as the Cisco Catalyst Series. Instead, it is a legacy reference to a WLAN controller (WLC) such as the Cisco 2006, 4400, ISR/WLCM, or Catalyst 6500 WiSM.)

  – Location servers
  – Location notifications

  Figure 8-55 presents a conceptual flow diagram that illustrates how WCS decides whether or not a notification is sent.

Figure 8-55 E-Mail Notification Logic Flow

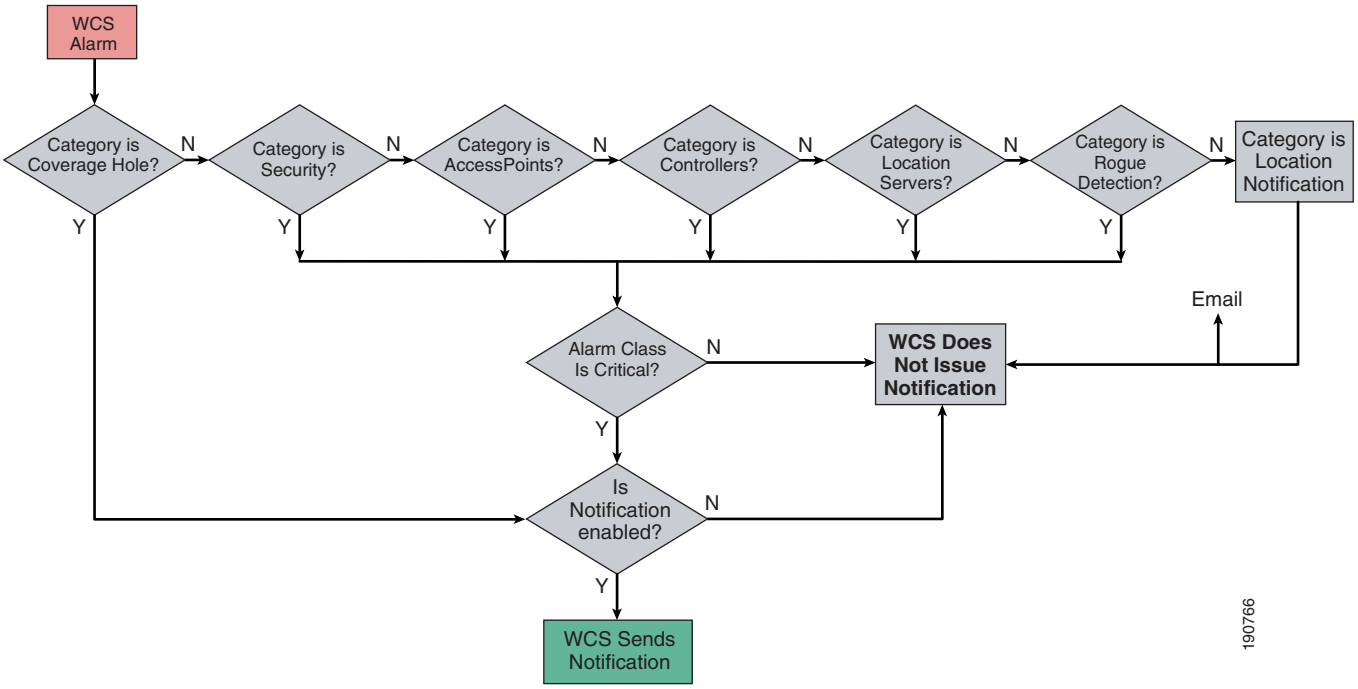

- If e-mail notification has been properly configured and enabled, the flow diagram indicates that WCS dispatches e-mails for triggered alarms if those alarms are either coverage holes alarms, or critical alarms. See Appendix B, "WCS Event and Alarm Severities." for a listing of alarms that are classified as critical alarms. Coverage hole alarms are classified as minor alarms and can also be found in Appendix B, "WCS Event and Alarm Severities.".

  The special case of the Location Notification alarm should be noted in Figure 8-55. Location Notification alarms are alarms that have been generated because of the reception of absence, containment, movement from marker, location changes, or battery level events from a location appliance. Because the location appliance itself is responsible for generating e-mail, syslog, SOAP/XML, and SNMP trap notification for these types of events, WCS does not provide for redundant notification capabilities and does not issue email notifications for Location Notification alarms.

Monitoring Alarms and Configuring E-mail Notifications, page 8-76 discusses how to configure WCS to successfully dispatch e-mail notifications.

Monitoring Events

As discussed previously, WCS maintains a log of all received events for a fixed seven-day retention period. These can be displayed by selecting **Monitor > Events** from the main menu selection bar. Because of the number of events that can be present in the event log in large networks, WCS provides the ability to filter the displayed events using the filter selector located in the left-hand margin of the **Monitor > Events** page. Using this tool, the displayed selection of events can be limited to a combination based on severity class and category.

Each column of the event display can be sorted by clicking on the column heading and choosing either ascending or descending sort sequence. Note that this can be very helpful not only in ensuring that the data you are viewing is in the proper date order but also in helping spot repeated patterns of events (that

is, sorting on message or failure type). There is only one hyperlink available in the event log data for further information and that is for the failure object. Clicking on any of the failure objects brings up additional information about the event in question.

Monitoring Alarms and Configuring E-mail Notifications

WCS maintains a log of all triggered alarms (shown in Figure 8-56) that can be displayed by selecting **Monitor > Alarms** from the main menu selection bar.

Figure 8-56 Monitor > Alarms

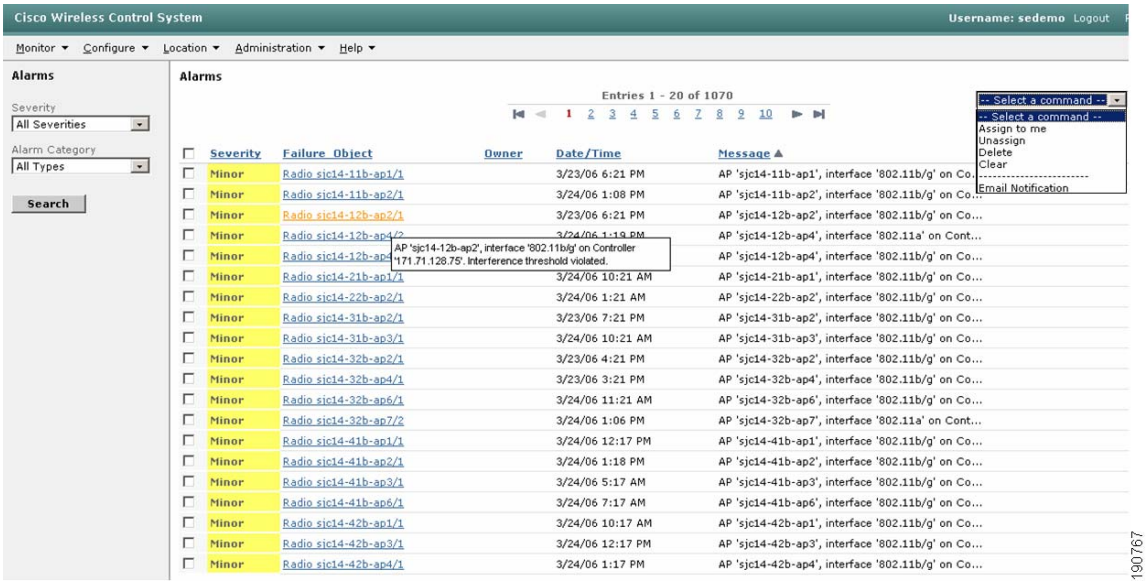

When performing a mouse-over of the failure object hyperlinks listed in the alarm display list, the alarm text associated with each message is displayed as shown in Figure 8-56. This mouse-over capability is useful when quickly scanning the list of alarms because it avoids the necessity of opening each alarm line item to simply determine the details. Each column of the alarm display list can be sorted by clicking on the column heading and choosing either ascending or descending sort sequence. This can be very helpful, not only in ensuring that the data you are viewing is in the proper date order but also in helping to spot repeated patterns in alarms.

As mentioned previously, alarms that have not been cleared (or manually deleted) remain in the database indefinitely. Cleared alarms remain in the database for a fixed period of 24 hours. In the case of very large networks, many alarms can be displayed; thus, WCS provides the ability to filter the displayed alarms using the selectors located in the left-hand margin of the alarms listing page. This allows the displayed selection of alarms to be limited by severity class and category, as shown in Figure 8-57.

Figure 8-57 Setting Display Filters in Monitor > Alarms

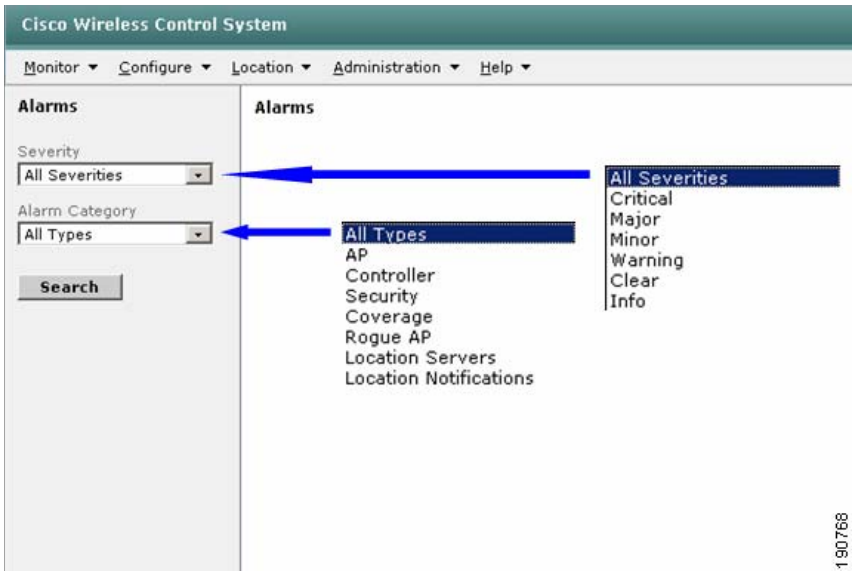

Depending on the alarm category, additional filtering options may be available. For example, when selecting the Rogue AP category, the filtering options shown in Figure 8-58 are available. Additional options become available when the AP Name, Floor Area, and Outdoor Area filters are selected to further qualify the rogue AP list.

Figure 8-58 Rogue AP Alarm Filtering Options

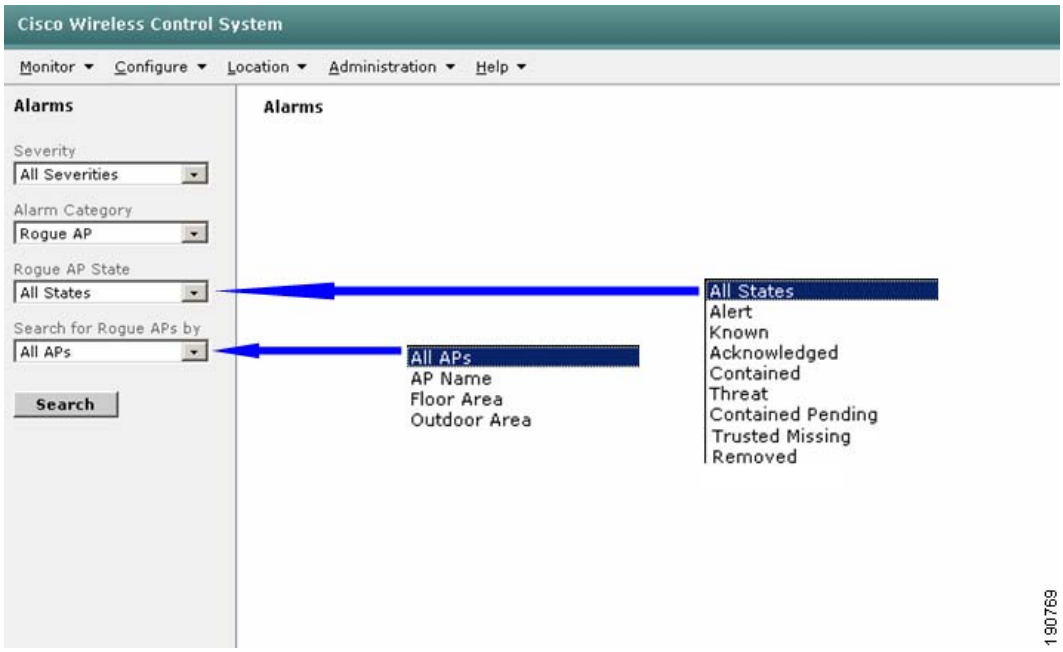

The options shown in Figure 8-58 for rogue AP state are defined as follows:

- Alert—Rogue access points that have been identified by the system as potential threats
- Known—Rogue access points identified as known internal rogues

- Acknowledged—Rogue access points identified as known external rogues
- Contained—Rogue access points that have been successfully contained by the system
- Threat—Rogue access points that are confirmed threats to the security of your network. An example of a confirmed threat is a rogue access point that has been identified by the system as physically attached to your internal wired network.
- Contained Pending—Rogue access points that are in the process of being contained by the system
- Trusted Missing—Known or acknowledged rogue access points that are no longer found
- Removed—Untrusted rogue access points that are no longer found

The drop-down command selector in the upper right-hand corner of the **Monitor > Alarms** alarms listing page provides the following options for most alarm categories:

- Assign to me—This command allows the WCS user to select alarms via their check boxes and to assign themselves as the "owner" of the alarm. Assigning owners to alarms is a useful administrative tool that can assist in managing alarm resolution by clearly indicating which person has agreed to take ownership of resolving the alarm. After you assign an alarm to yourself, the user name with which you are currently logged into WCS is displayed as the owner in the **Monitor > Alarms** display as well as in the alarm detail display. Alarms can be assigned only to the user name that you used when logging into WCS; you may not assign alarms to other users.
- Unassign—A complement to "Assign to me", this option allows you to remove the owner of an alarm.
- Delete—Removes the alarm from the alarm database entirely (that is, as if the alarm never existed). Note that although this command removes the alarm, the underlying events that caused the alarm are not removed. In the rare circumstance of an entry persisting in the database even after a manual clearing of the alarm has been performed, the delete command can be used to forcibly remove the alarm. The deletion of an alarm is not reversible except via restoration of the WCS database. Use of the delete command can be restricted via permissions assigned in **Administration > Accounts > Groups**.
- Clear—Manually clears a currently active alarm by issuing a clear alarm that replaces the original alarm severity in much the same manner as a clear alarm received from a network component clears the corresponding active alarm. After an alarm is cleared, only the clear alarm remains listed in **Monitor > Alarms** (the original alarm severity is replaced by the clear alarm). 24 hours after the alarm is cleared, the clear alarm is removed from the database. Note that clearing an alarming does not affect the underlying events that caused the alarm (that is, the events are still present in the event log for seven days).
- E-mail notification—Opens the page shown in Figure 8-59, and is where the communication parameters are configured for e-mail notification. Note that as mentioned in Relationship Between Traps, Events, Alarms, and Notifications, page 8-69, there are seven categories of alarms and each category allows you to configure different e-mail settings.

Figure 8-59 Monitor > Alarms > E-mail Notification

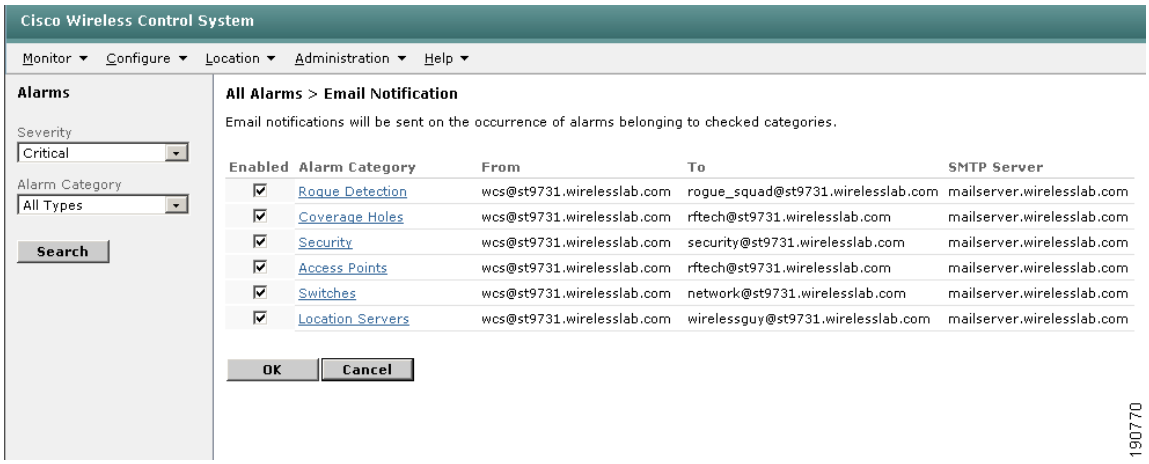

To configure the parameters for one of the seven categories, click on the hyperlink for the alarm category. This displays the page shown in Figure 8-60.

Figure 8-60 Specifying E-mail Parameters

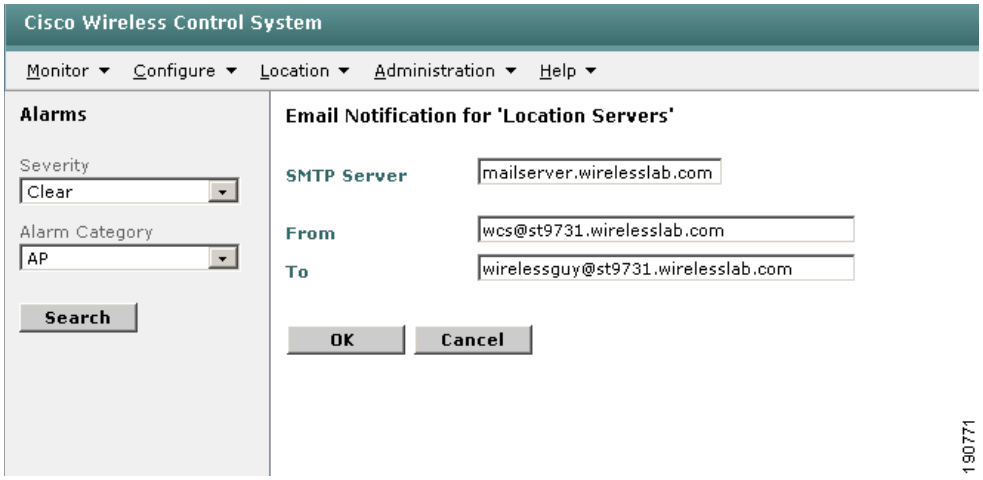

Within each category, only one SMTP server can be specified (either as a fully qualified domain name or IP address); however, multiple e-mail destination addresses can be specified and separated by commas. When specifying multiple e-mail destinations, be aware that there is a 56-byte total length limitation on the “To” field.

Notice that there is no e-mail notification configurable for the Location Notifications alarm category, as described in Relationship Between Traps, Events, Alarms, and Notifications, page 8-69 and Figure 8-55.

When you have specified the e-mail notification parameters, click **OK** and return to the e-mail notifications screen. Before your e-mail notifications become active, you must enable them as shown in Figure 8-59 by checking the applicable check box(es) in the Enabled column and clicking **OK**. Your e-mail notification configuration is complete at that point and e-mail notification should be functional.

An example of an actual e-mail alert received can be seen in Figure 8-61.

Figure 8-61 E-mail Notification

Date: Tue, 21 Feb 2006 12:37:29 -0500 (EST)
From: wcs@st9731.wirelesslab.com
To: rftech@st9731.wirelesslab.com
Subject: Access Points Alarm from Radio AP1000#3/2

TIME:Tue Feb 21 12:37:29 EST 2006

An Alert of Category AP is generated with severity 1
by Radio AP1000#3/2 .
The message of the alert is AP 'AP1000#3', interface '802.11a' is down on Controller '10.1.56.18'.

190772

Compared to the other alarms categories available under **Monitor > Alarms**, rogue AP alarms are somewhat of a special case in that the options presented under the drop-down command selector are expanded, as shown in Figure 8-62.

Figure 8-62 Monitor > Alarm Rogue AP Alarms

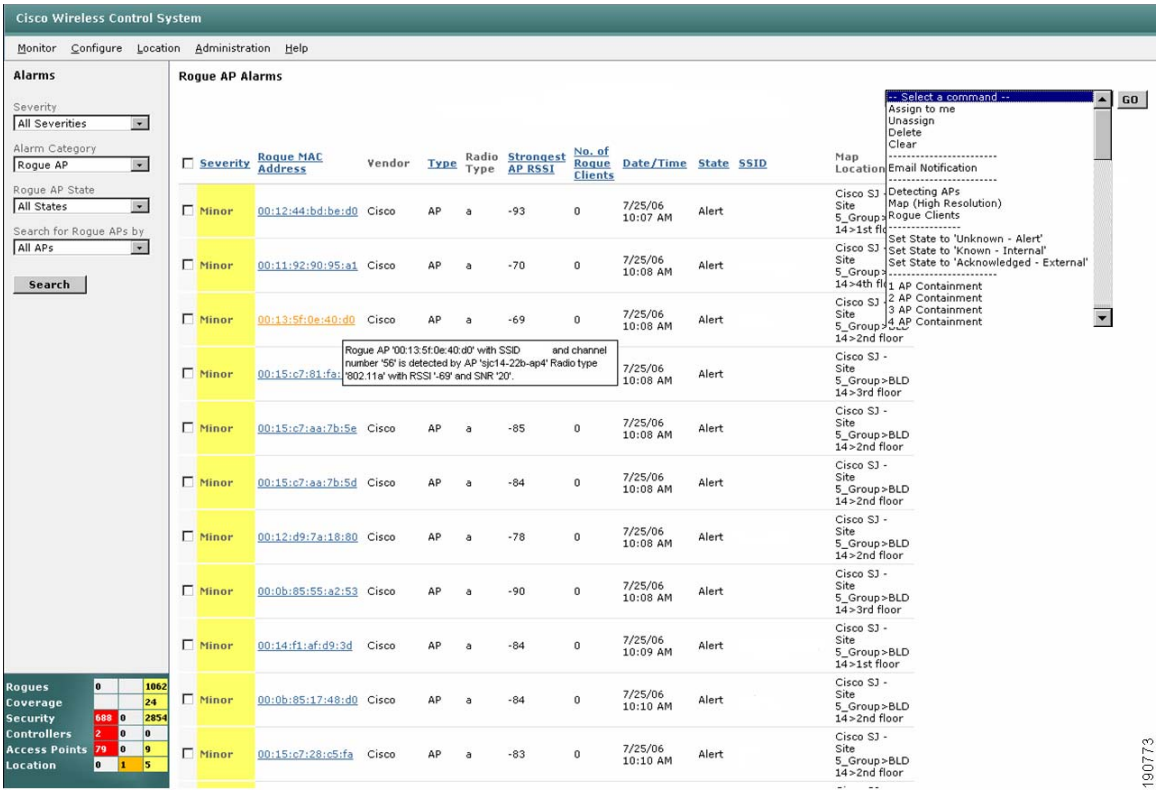

190773

The following additional command options are available under rogue AP alarms:

- Detecting APs—Provides a listing of all lightweight infrastructure access points detecting the selected rogue AP. When using this option, you must select only one rogue access point.
- Map—Provides either a low-resolution or high-resolution on-demand location display of the current location of the rogue AP. See On-Demand Location of Individual Rogue Access Points, page 8-87 for further details about on-demand rogue access point location.

- Rogue Clients—Provides a listing of the rogue clients that are associated to this rogue access point. See Monitoring Security, page 8-65 for more information, and Figure 8-51 for an example of the rogue client detail screen.
- Set State to “Known–Internal”—When searching for rogue APs, this state is referred to as “Known”.
- Set State to “Acknowledged–External”—When searching for rogue APs, this state is referred to as “Acknowledged”.
- AP Containment—Initiates rogue AP containment using from one to four infrastructure access points.

Clicking on the failure object hyperlink of any of the items listed on the **Monitor > Alarms** listing page results in the Alarms Detail page being displayed. The look and feel of the alarm detail page varies somewhat based on the type of alarm being displayed. A common alarm detail format shared among most alarms (except for the rogue AP alarm) is shown in Figure 8-63.

Figure 8-63 Alarm Detail Page—Annotations

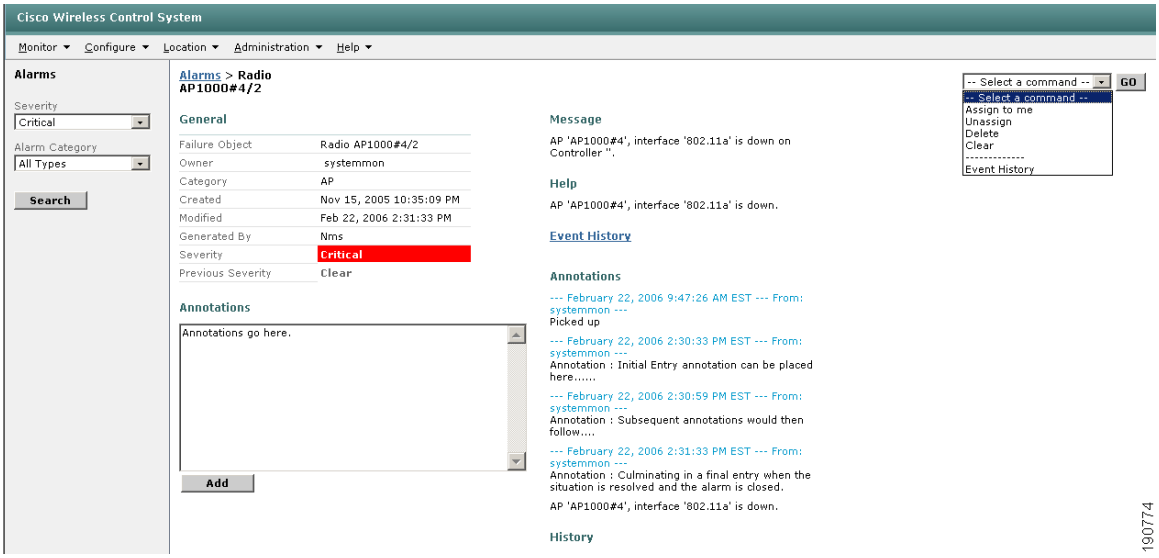

In all cases, the value shown for the “Generated By” field in the alarm detail page shown in Figure 8-63 indicates the source of the information that triggered the alarm or event:

- Device—Indicates that the alarm or event was generated based on information obtained from an SNMP trap received from the device.
- NMS—Indicates that the alarm or event was generated based on information obtained during SNMP polling.

Note the use of the Annotations area for keeping a running log of what has occurred and who has been involved in resolving the alarm. Annotations are added via the entry box on the left-hand side and appear in the order in which they were added in the annotations area on the right. The date and time of the alarm assignment is indicated by the “picked up” entry (if the alarm were to be unassigned, an “unpicked” entry would show up here as well). The clearing of an alarm results in a severity change but no additional entry in the annotations area.

Rogue AP alarms use a somewhat different detail page format, with additional information provided about location, location notification, and any rogue clients that might be associated to this rogue AP available. The rogue AP alarm detail page is shown in Figure 8-64.

Figure 8-64 Rogue AP Alarm Detail

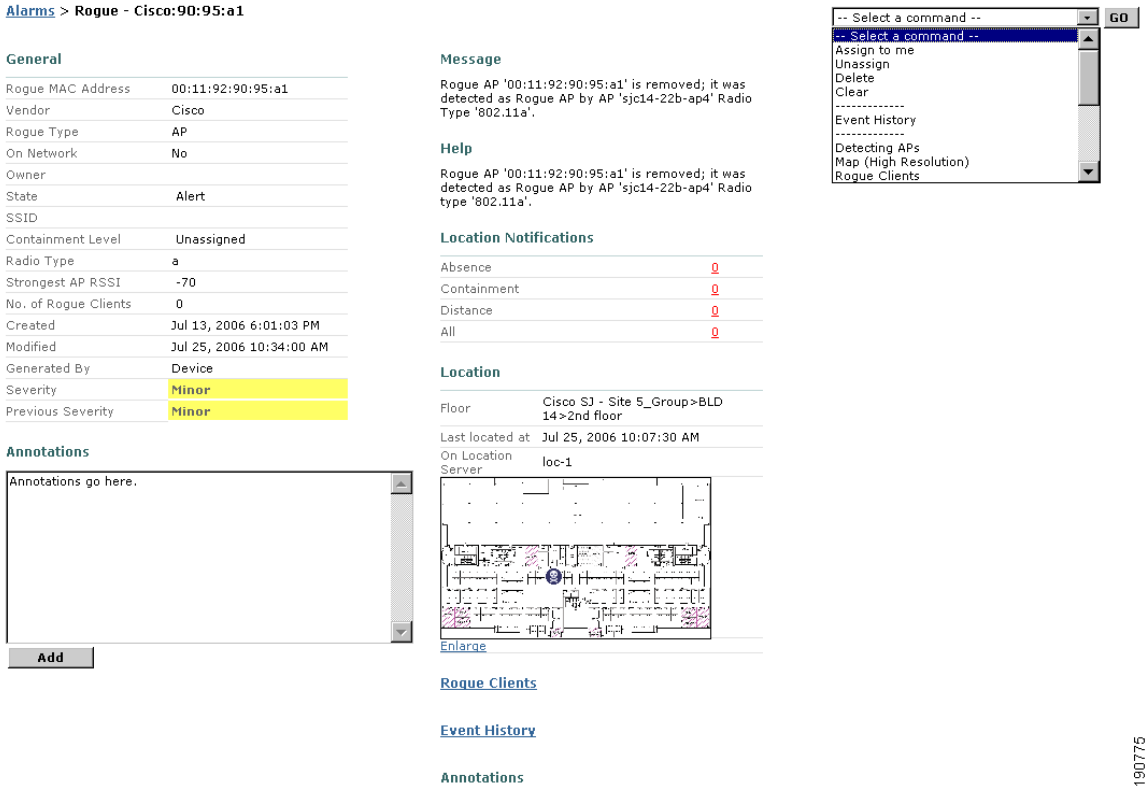

The drop-down menu selector in the upper right-hand corner of the rogue AP alarms page offers an expanded set of command options when compared to other alarms. The majority of the command options are familiar from the discussion of rogue AP alarms listings. Two options found here that are not available when listing rogue AP alarms are event history and location history. Event history is available both as a command drop-down selection and as a hyperlink on the rogue AP alarm detail page. The function in both cases is the same; the list of events that are associated with this alarm are displayed. Location history is available for the rogue access point in a similar fashion to that described for WLAN clients in Monitoring Clients, page 8-54.

Using WCS to Locate Devices in Your Wireless Network

WCS offers the ability to locate rogue client devices, 802.11 active RFID tags, rogue access points, and rogue clients when they are detected within your wireless network. Location of these devices can be provided by WCS on either an on-demand basis for a single device or on a routinely updated basis for multiple devices when used with a location appliance.

For users requiring location services only when the need arises to locate a lost device, the on-demand location capabilities afforded by WCS-Base or WCS-Location may be all that is required. However, for users that rely on the ability to track the movement of devices in their environment on a regular basis, require alarms and notifications when devices move into or out of defined areas, maintain location history for more than seven days, or interface to third-party location client applications, the use of a location-licensed WCS server with the location appliance is a more appropriate choice.

This section describes each of these options in further detail.

On-Demand Device Location

In this section, "on-demand" device location refers to the ability of WCS to display the current position of a single device whenever you explicitly request the system to do so without the use of a location appliance. Although both the base as well as the location-enabled version of WCS can perform on-demand location, they vary both in for which devices each provide positioning information as well as the level of position granularity.

Only a WCS server that is licensed for base-level functions can perform on-demand client and rogue device location, by placing an icon on a floor map nearest the access point that has detected the device with the highest signal strength. Knowing which infrastructure access point detects the device with the highest signal strength usually provides sufficient resolution for casual location services use, especially in cases where the need for location services in a particular business situation may not justify additional investment in software or hardware.

A location-licensed WCS server improves these on-demand capabilities by providing "high-resolution" location of WLAN clients and rogues using Cisco RF Fingerprinting positioning technology. As opposed to simply knowing the lightweight access point that detected the client with the highest signal strength, RF Fingerprinting can provide location accuracy of 10 meters, 90 percent of the time (90 percent precision) in a properly-designed system. A location-licensed WCS server can perform on-demand location of a single device at a time, and is an excellent choice where higher accuracy is desired to reduce the amount of time and effort that must be expended searching for items that are within the range of a particular access point.

The location appliance does not play an active role in establishing on-demand location because it is typically entirely driven using the client and rogue information available in WCS databases and the WLAN controllers. Rather, the Wireless Location Appliance allows WCS to display the location of multiple devices simultaneously by performing location calculations on the data acquired by the location appliance during its polling of WLAN controllers. The location appliance polls based on configured polling parameters and does not poll WLAN controllers on-demand. The information that is used to establish device location during an on-demand location request is supplied by WCS and the WLC.

On-Demand Location of WLAN Clients

On-demand location of a single WLAN client at a time without the use of a location appliance can be performed from WCS using the **Monitor > Devices > Clients** Summary menu shown in Figure 8-39. To do so, follow these steps:

Step 1 Using the "Search for Clients by" feature in the left-hand margin of the page narrow the search to include only the clients of potential interest. Be sure to specify "WCS Controllers" instead of "Location Servers" if there is not a location appliance present. Click **Search**.

Step 2 Click on the client user name hyperlink of the client for which you want to display location.

WCS displays the client detail screen shown in Figure 8-65 for the client you have selected.

Figure 8-65 Client Details

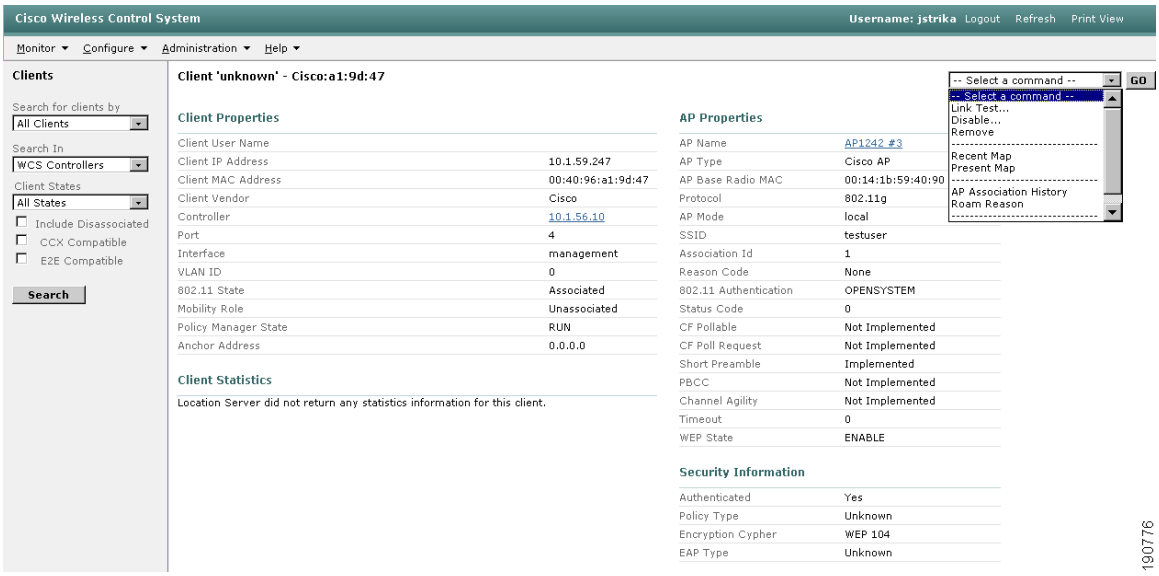

Step 3 Click on the command drop-down menu in the upper right-hand corner of the screen and choose from either the “Recent Map” or “Present Map” options.

- Recent map—Displays the last recorded historical location of the client from the WCS database and is non-intrusive to the WLAN client. The information that is used to establish device location in this case is garnered via the routine background polling cycles of WLAN controllers by WCS.
- Present Map—Displays the current location of the client and can cause disruption to the existing client session. To gather the current signal strength information of the client, the client is de-authenticated/disassociated very briefly and must then re-associate/re-authenticate. The amount of time this takes depends on the details of the client authentication method being used. Although this disruption is typically quickly recovered, some applications (such as voice and some business data applications) may prove to be more sensitive to the interruption than others. For this reason, Recent Map instead of Present Map is generally preferred when performing on-demand location for active, in-session users unless the impact of any such interruption has been assessed beforehand.

  In most cases, if on-demand location is being used to locate a lost device, this interruption has little impact because if the device is still powered on, it is typically not in use. However, lost devices tend to eventually power down because of battery exhaustion, and in that case the Recent Map option is of much more use in determining the last known location.

Choose either of the mapping options and then click **Go**.

Step 4 WCS systems that are licensed only for base-level functionality (WCS-Base) display on-demand location of clients in a manner similar to what is shown in Figure 8-66. Note that the icon for a WLAN client is placed adjacent to the access point that detects the client with the highest signal level. The location of the client on the map does not indicate the estimated location of the client, just that it has been detected with the highest signal strength by the infrastructure access point that it is adjacent to. In Figure 8-66, that access point is AP1000#3.

Figure 8-66 On-Demand WLAN Client Location (Base Level WCS)

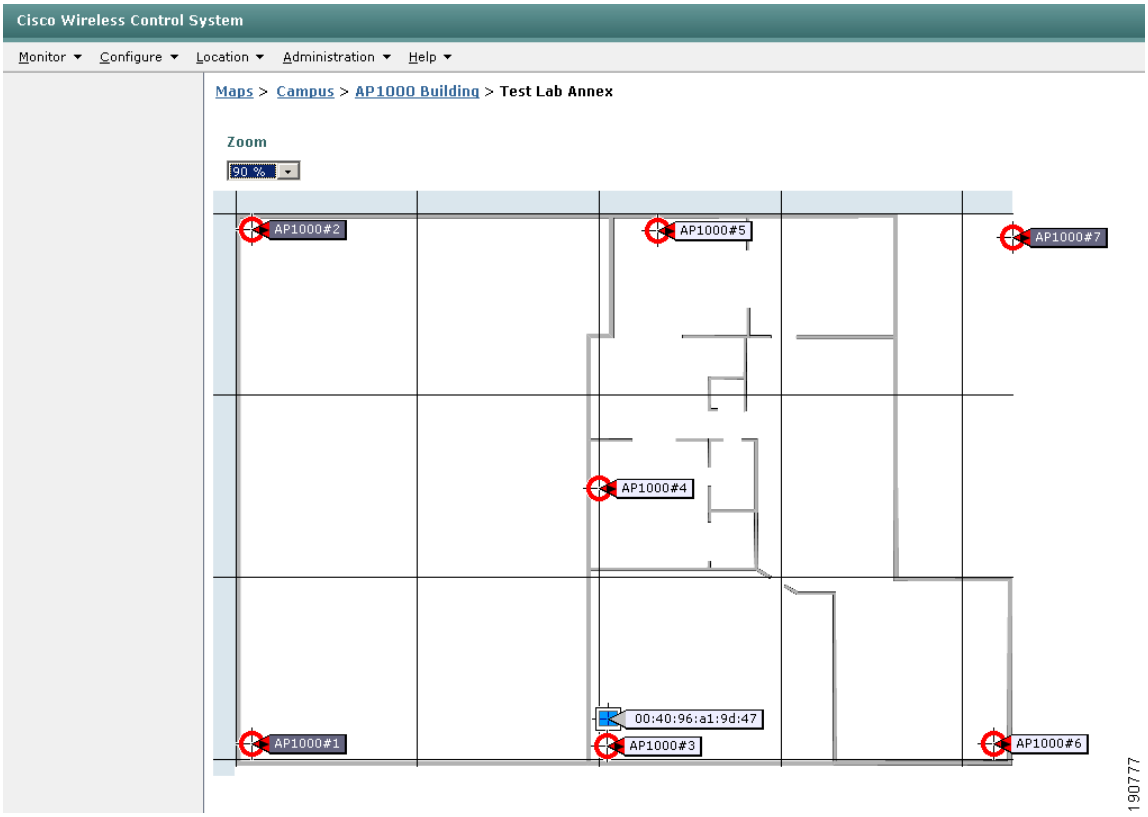

When using a location-licensed WCS server without a location appliance, the process is essentially the same with the exception that WCS uses RF Fingerprinting to derive estimated client location on a "high-resolution" map. In this case, your on-demand location results resemble Figure 8-67 instead of Figure 8-66.

Figure 8-67 High Resolution On-Demand Client Location (Location-Licensed WCS)

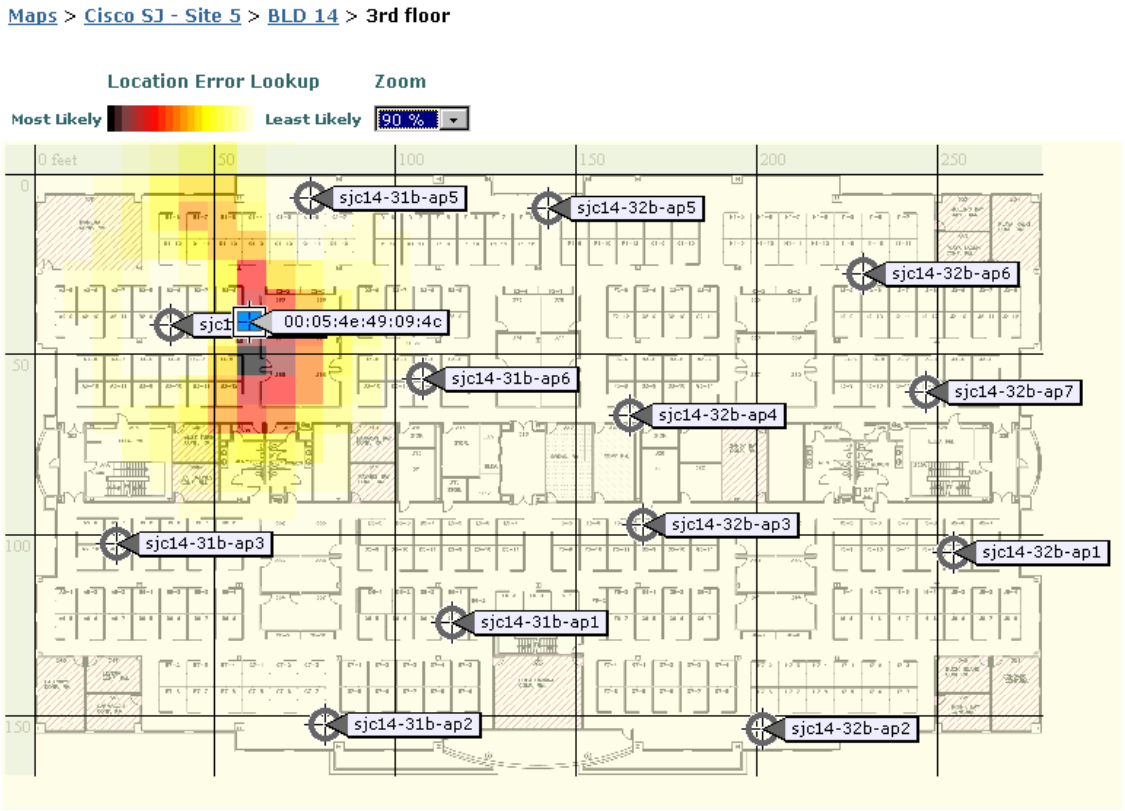

The main difference between what Figure 8-66 and Figure 8-67 is that in Figure 8-67, the placement of the blue rectangular WLAN client icon is intended to represent the estimated client position on the map. The various color bands in the display indicate varying location probabilities, as shown in the "Location Error Lookup" legend at the top of the display. Performing a mouse-over of the various colors in the legend itself displays the error probability associated with each color band. In contrast, the positioning of the blue rectangular WLAN client icon in Figure 8-66 is not intended to convey estimated client position on the map, but merely to indicate which access point has detected the WLAN with the greatest signal strength.

On-Demand Location of Individual 802.11 Active RFID Asset Tags

The 802.11 active RFID asset tags that can be tracked by the Cisco Location-Based Services solution can be grouped into two basic categories:

- 802.11 Active RFID asset tags that communicate via Layer 2 (L2) multicasts such as the AeroScout T2 asset tag. These asset tags typically use the WDS frame format and do not associate to the WLAN infrastructure. These asset tags appear as yellow tag icons within WCS floor maps.

Note For a complete discussion of the AeroScout T2 tag and WDS frame formats, see *Wi-Fi Location Based Services—Design and Deployment Considerations* at the following URL: http://www.cisco.com.

- 802.11 Active RFID asset tags that communicate as full WLAN clients and associate/authenticate to the WLAN infrastructure, such as PanGo Locator LAN tags. These types of asset tags are viewed as WLAN clients by WCS and the location appliance, and they appear as blue rectangular icons within WCS floor maps.

To locate 802.11 Active RFID asset tags that are of the latter category, see On-Demand Location of WLAN Clients, page 8-83 because these types of asset tags are technically treated as WLAN clients by the Cisco UWN.

Note that beginning with release 4.0 of WCS, a location appliance is required to determine the position of any Layer 2 multicast-based RFID asset tags. A subsequent maintenance release is expected to reinstate the capability to perform on-demand location of individual Layer 2 multicast-based RFID tags when using a WCS server that is licensed for location use. In addition, beginning with release 4.0, WCS servers that are licensed only for base-level functionality do *not* have the capability to perform *any* type of location determination for Layer 2 multicast-based RFID tags.

On-Demand Location of Individual Rogue Access Points

On-demand location of a single rogue access point at a time without the use of a location appliance can be performed with base-level WCS using the Rogue AP Alarms menu accessible via **Monitor > Security > Rogue APs**. To do so, select a single Rogue AP by enabling its check box. From the command drop-down located at the upper right-hand corner of the screen, select "Map" and then click **GO**.

When using the base-level WCS, the screen shown in Figure 8-68 is displayed.

Figure 8-68 On-Demand Rogue Access Point Location (Base-Level WCS)

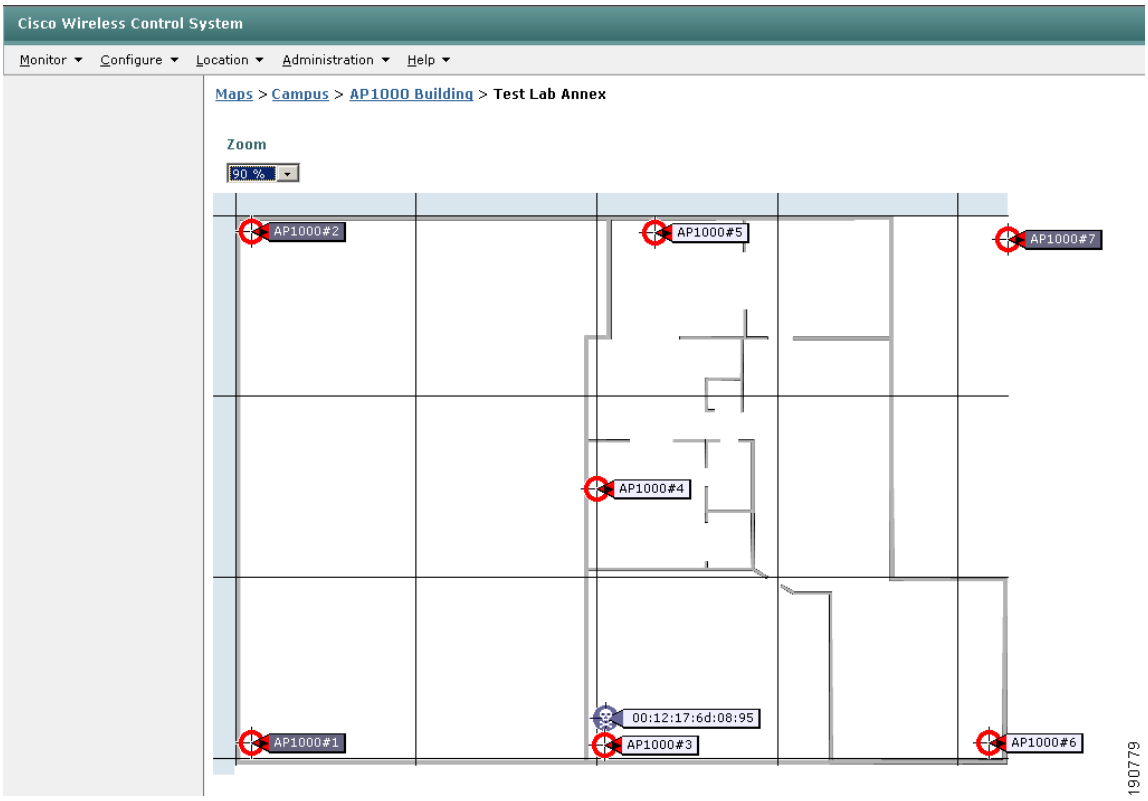

Note that the circular black "skull-and-crossbones" icon representing a rogue access point is placed directly adjacent to the lightweight infrastructure access point that has detected it with the highest signal level.

When using a version of WCS that is licensed for high resolution location tracking, the process is identical with the exception that WCS uses RF Fingerprinting to determine the estimated position of the rogue access point and displays it on a "high resolution" map similar to that shown in Figure 8-67. With the exception of the icon used to indicate the location of the rogue AP, the look and feel of the high-resolution on-demand map is the same as that shown in Figure 8-67.

On-Demand Location of Individual Rogue Clients

On-demand location of a single rogue client at a time can be performed from WCS without a location appliance using the **Monitor > Security > Clients** Summary menu shown in Figure 8-39. To do so, follow these steps:

Step 1 Using the "Search for Clients by" feature in the left-hand portion of the screen, specify information to narrow the search to include only rogue clients of potential interest. Be sure to specify "WCS Controllers" instead of "Location Servers" if there is not a location appliance present.

Click **Search**.

Step 2 WCS displays a listing of detected rogue clients that are found in the WCS database. Click on the MAC address hyperlink of the rogue client for which you want to display location.

Step 3 WCS displays detailed information about the rogue client including the time it was detected, which access points detected it, and at what signal levels. Click on the command drop-down menu in the upper right-hand corner of the screen, select "Map", and click **GO**.

The base-level WCS displays a location map with a black rectangular icon representing a rogue client, as shown in Figure 8-69. It is placed nearest the infrastructure lightweight access point that has detected it with the highest signal level.

Figure 8-69 On-Demand Rogue Client Location (Base Level WCS)

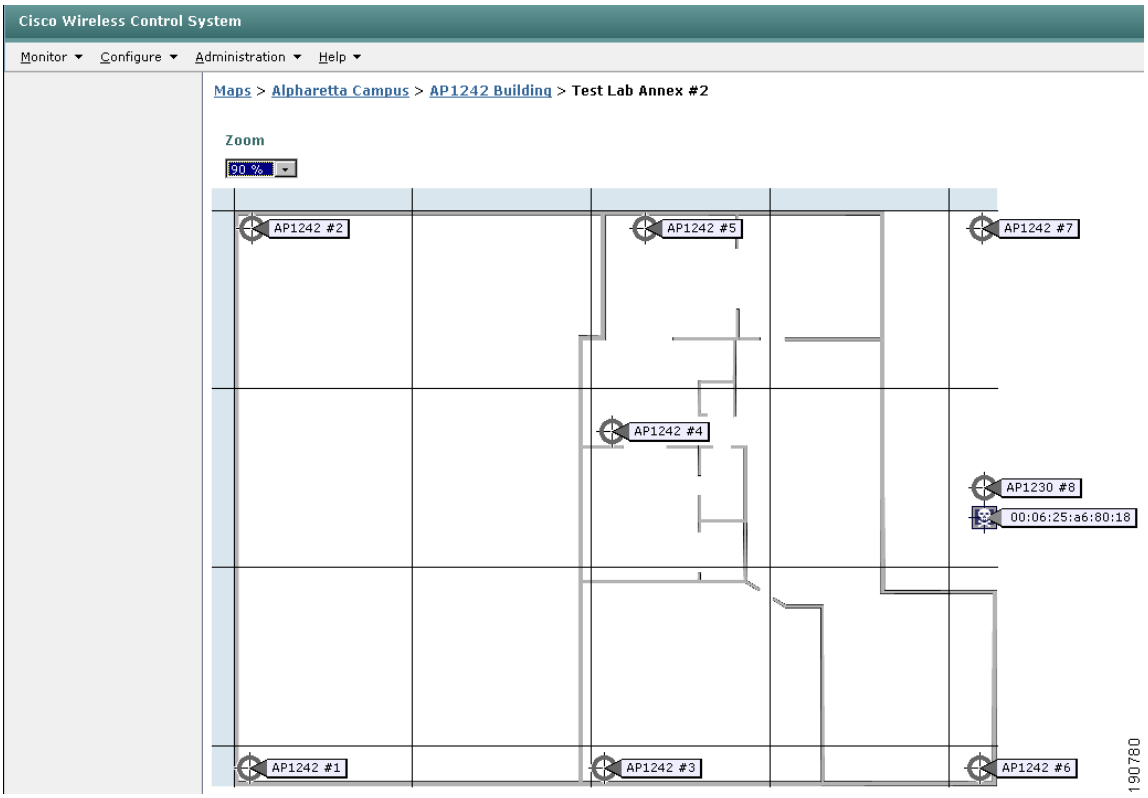

When using a version of WCS that is licensed for high resolution location tracking, the process is essentially the same with the exception that WCS uses RF Fingerprinting to determine the estimated position of the rogue client and displays it on a "high resolution" map similar to that shown in Figure 8-67. With the exception of the icon used to indicate the location of the rogue client, the look and feel of the high-resolution on-demand map is the same.

WCS and the Location Appliance

When a Cisco Wireless Location Appliance is added to a location-licensed WCS server, its high-resolution location capabilities are enhanced by the ability of the location appliance to issue location notifications and compute positioning information for multiple devices simultaneously while maintaining a much larger amount of location history data in its internal databases.

The location appliance interfaces to WCS using the SOAP/XML API interface. Via this API, WCS serves in the role of both a *location client* in displaying the location of many simultaneous devices, asset tags, and rogues on location maps as well as a *control client* in acting as the primary user interface to the location appliance and handling its configuration.

This same SOAP/XML API allows for integration with other business applications that can use the location information contained within the location appliance for a variety of creative applications. Asset tracking, inventory management, location-based security, and automated workflow management are just a few examples of this. Third-party location client applications access the location appliance only via the API and typically do not access WLAN controllers or other components of the network directly.

Operators can configure location appliances to collect client RSSI data and statistics from WLAN controllers at defined intervals. The location appliance allows WCS to display the location of multiple devices simultaneously by performing location calculations on the data acquired by the location appliance during its polling of WLAN controllers. The location appliance polls based on configured polling parameters and does not participate in on-demand location display.

The location appliance also provides *location-based event notification*, whereupon it generates e-mail and other notifications directly to specified destinations. These alarms and notifications can be triggered through area boundary, allowed areas, and distance definitions in the location appliance. These alarms and notifications can also provide advanced warning of rogue movement and appearance/disappearance. Using WCS, you can configure location appliance event notification parameters that allow the location appliance to send notifications to destinations configured via the WCS **Location > Notifications** menu option. The location appliance can be defined to transmit messages using SOAP, SMTP, SNMP traps, or syslog messages if clients or assets become missing, enter or leave coverage areas, or stray beyond a set distance from a pre-determined marker.

Architecture Overview

The overall architecture of the Cisco location-based services solution is shown in Figure 8-70.

Figure 8-70 Cisco Location-Based Services Architecture

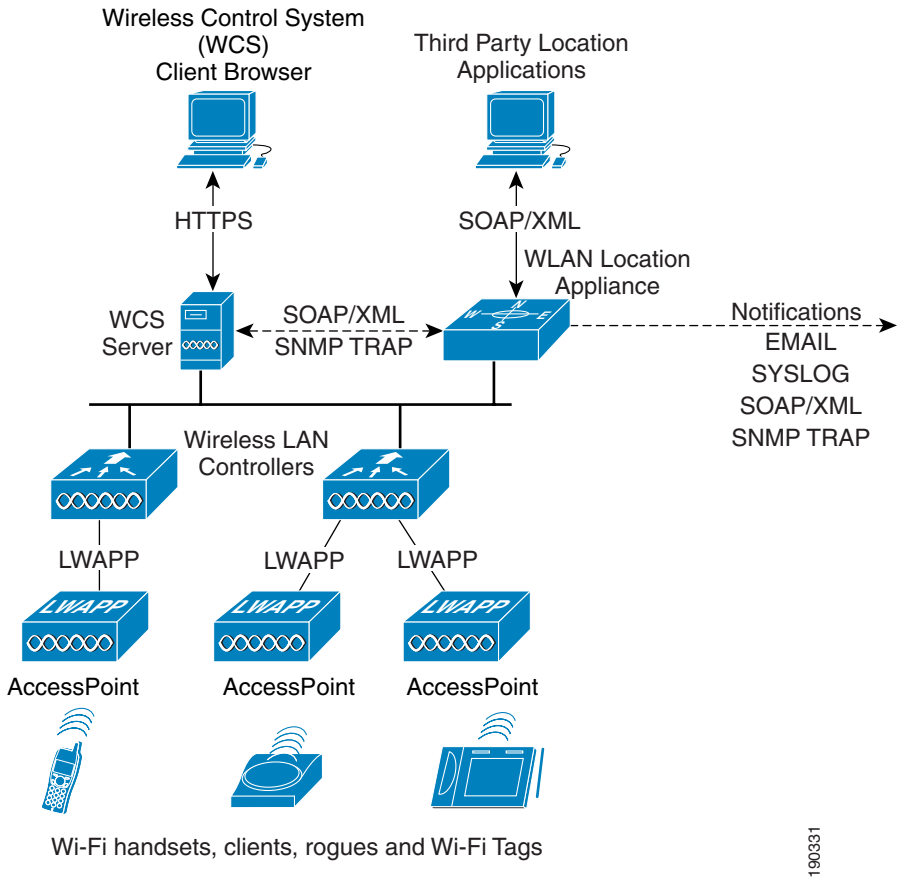

Access points forward information to WLAN controllers about the detected signal strength of any Wi-Fi clients, 802.11 active RFID tags, rogue access points, or rogue clients. In normal operation, access points collect this information on their primary channel of operation, going off-channel and scanning all

channels in their regulatory channel set periodically. The collected information is forwarded to the WLAN controller with which the access point is currently registered. Each controller manages and aggregates all such signal strength information coming from its access points. The location appliance uses SNMP to poll each controller for the latest information for each tracked category of device. In the case of a location tracking system deployed without a location appliance, WCS obtains this information from each controller directly.

WCS and the location appliance exchange information about calibration maps and network designs during a process known as *synchronization*. During a *network design synchronization* between WCS and the location appliance, the up-to-date partner updates the design and calibration information of the out-of-date partner. The location appliance synchronizes with each controller containing access points participating in location tracking during *controller synchronization*. The synchronization of notification schedules and destinations between the location appliance and WCS is handled via a process referred to as *event group synchronization*. Synchronization occurs either on-demand or as a scheduled task, the timing of which is determined by the **Administration > Scheduled Tasks** main menu option in WCS.

Location information is displayed to the end user using a *location client* application in conjunction with the Cisco Wireless Location Appliance. Typically, this role is fulfilled by the Cisco WCS, which is capable of displaying a wide range of information about the location of clients, asset tags, rogue access points, and rogue clients. However, location client functionality is not limited to WCS, because other third-party applications written in accordance with the Cisco Location Appliance API and using the SOAP/XML protocol can also serve as a location client to the Wireless Location Appliance (as shown in Figure 8-70).

Although it is most common to have a single WCS paired with a single location appliance, variations on this theme are also possible to better support specific configurations. For example, in some cases because of a high number of clients and rogues, the supported device tracking capability of a single location appliance may be exceeded while being substantially under the supported access point and controller limitations of WCS. In this case, you can partition the network from the perspective of location service only and assign one half of the network (for example) to location appliance A, with the other half to location appliance B. Both of these location servers can then be managed under a single common WCS. Using this approach, the WCS administrator is presented with the simplicity of managing a single management domain from a single WCS server, but with the extensibility of multiple location domains that allow for addressing client, asset tag, and rogue counts that exceed the supported capacities of a single location appliance.

Tracking Clients, Asset Tags, and Rogues with the Location Appliance

As mentioned in On-Demand Device Location, page 8-83, WCS can perform on-demand location of clients, rogue access points, and rogue clients regardless of the presence of a location appliance. However, the addition of the location appliance allows you to view the location of multiple devices across all these classes simultaneously and track their history for longer periods than WCS allows.

Figure 8-71 provides a visual example of this multiple device simultaneous tracking capability. It may be helpful to compare the capabilities shown here to those shown in Figure 8-66 and Figure 8-67 for on-demand location. In Figure 8-71, devices from all supported categories are displayed on a floor with a periodic screen refresh set to five minutes. The floor plan loaded by the WCS administrator used as a background onto which is superimposed an RF “heatmap” showing the predicted RF coverage of each access point. Access points are shown on the floorplan along with internally generated channel information tags. Using information contained in the location appliance database, WCS displays icons for each device at locations that were determined by the location appliance using RF Fingerprinting positioning calculations.

Figure 8-71 Floor Map—Simultaneous Tracking of all Device Categories

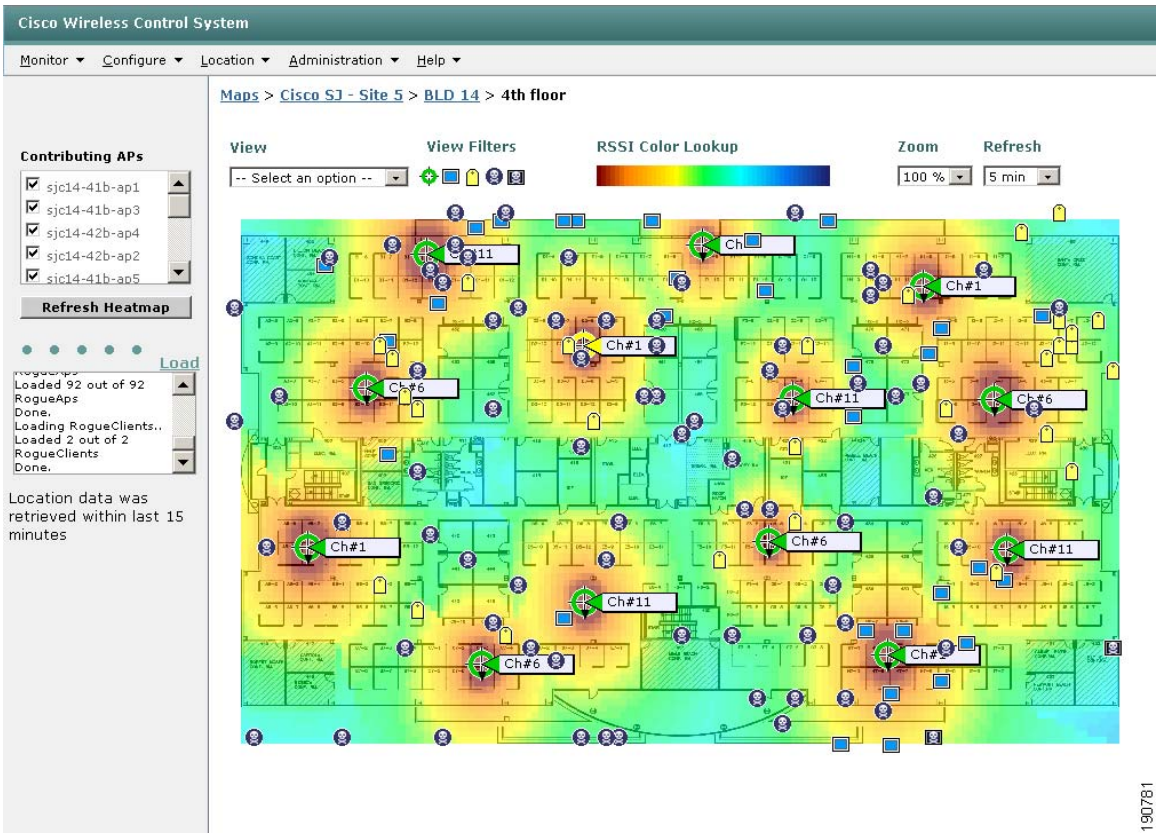

For a comprehensive discussion of the device tracking capabilities that are available when using WCS with the Wireless Location Appliance, see *Wi-Fi Location Based Services—Design and Deployment Considerations* at the following URL: http://www.cisco.com.

Using WCS to Efficiently Deploy Your Wireless Network

One of the challenges facing network designers and installers is how to efficiently and rapidly deploy large numbers of lightweight access points and controllers across numerous sites.

In very large deployments, the ability to rapidly deploy and configure the network infrastructure is a key aspect of a successful and economically successful project implementation. It is not always feasible to deploy experienced network technicians to perform configuration, troubleshooting, and software upgrades on-site to each and every site in a large-scale deployment. Often there are simply too many sites being installed simultaneously and not enough personnel available to make this practical.

To address this situation, the wireless networking solution should require only minimal, basic configuration on-site at installation time to allow IP connectivity back to a centralized management site. This allows further configuration to be performed over the network by a central pool of experienced technical personnel. Experienced central site technical staff can make use of configuration efficiencies present within the WCS management system that greatly reduce the number of steps required to configure WLAN controllers and lightweight access points in accordance with all policies and standards set forth for the enterprise.

In the Cisco Unified Wireless Network, each controller requires only basic interface and IP configuration before it is accessible over an IP network. This is typically performed via the controller serial console and is a relatively simple affair. After interfaces are configured and the controller is attached to the IP network, WCS can be used to complete the application of configuration parameters to WLAN controllers as well as the latest software levels. After configuration of WLAN controllers has been completed, there is no need to individually configure lightweight access points because they derive their configuration parameters as well as their internal operating software in a "zero-touch" fashion from the WLAN controllers themselves.

Several WCS features are described in this chapter that can be particularly helpful during the deployment of large wireless networks across numerous sites, allowing far more efficient configuration than would be possible by simply accessing each controller using a web browser or CLI session and performing all configuration manually.

Policy Templates

Policy templates (Defining and Applying Policy Templates, page 8-22) are a key feature of WCS that can reduce the amount of effort required to configure remote site WLAN controllers during a deployment of any size. Policy templates allow for a set of related configuration objects to be defined, applied to selected controllers, and then saved for later use with subsequent controllers awaiting deployment. As each controller is physically installed and made available over the IP network, policy templates can be applied to one or more controllers, access points, or radios by the WCS operator. (Access points and radios must be registered with controllers to be eligible for configuration via templates. See Defining and Applying Policy Templates, page 8-22 for further information.)

When configuring multiple WLAN controllers that are part of the same mobility group, consider using the Configuration Groups facility described in Using Policy Template Configuration Groups, page 8-25 as opposed to applying templates one at a time. The use of configuration groups allows you to assign multiple controllers to a mobility group and apply one or more templates to them, saving a considerable amount of labor.

An example of how policy templates can be used to save both time and effort can be seen in the following steps taken to remotely configure a newly-installed Cisco 4400 WLAN Controller at a remote site.

The WLAN controller to be installed should receive its initial basic configuration via a local CLI session using the serial port. This is typically performed either before shipment at a staging center or at the remote site during installation by on-site installers. Appendix C, "Example of Wireless LAN Controller Initial Setup," indicates an example of the controller Setup Wizard and the type of information that is required for basic setup. Complete guidance and step-by-step instructions for configuring a WLAN controller using the controller Setup Wizard can be found in the document entitled Cisco 4400 Wireless LAN Controller—Quick Start Guide.

Step 1 After the controller has received basic configuration and it has been confirmed that the controller has been successfully attached to the network, it attempts to add the controller to WCS as described in Adding Controllers, page 8-8. If this fails, re-check connectivity and SNMP parameters specified in **Configure > Controllers > Add Controller**.

Step 2 Apply all desired templates to the WLAN controller as described in Defining and Applying Policy Templates, page 8-22 and *Cisco Wireless Control System Configuration Guide, Release 4.0* at the following URL: http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_book09186a00806b57ec.html.

Step 3 Save the controller configuration to nonvolatile (flash) memory as described in Configuring WLAN Controllers, page 8-12 and Figure 8-10.

Step 4 If necessary, upgrade any controller software as described in Managing Controller Operating Software, Web Authentication Bundles, and IDS Signatures, page 8-28. If a controller operating system upgrade was performed, reboot the controller by visiting the All Controllers screen shown in Figure 8-3 and selecting the controller you wish to reboot.

Step 5 Select **Reboot Controllers** from the upper right-hand drop-down menu and click **GO**. After the controller has fully rebooted, verify that it contains the expected software version by viewing **Monitor > Devices > Controllers > Summary** and observing the software version listed under Inventory.

Step 6 Have the lightweight access points at the remote site connected to the local network and allow them to boot up completely as per the detailed guidelines found in either Quick Start Guide LWAPP-Enabled Cisco Aironet Access Points or Cisco Aironet 1240AG Series Lightweight Access Point Hardware Installation Guide, which are both available at the following URL: http://www.cisco.com.

Step 7 After the access points have successfully booted up and can be seen to have registered to the controller as per Monitoring Access Points, page 8-51 and Figure 8-36, proceed to define the AP and radio policy templates for this controller and apply them to the access points as described in Defining and Applying Policy Templates, page 8-22 and *Cisco Wireless Control System Configuration Guide, Release 4.0* at the following URL: http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_book09186a00806b57ec.html.

If the policy templates are successfully applied to the access points and radios, they are automatically saved within those devices and retained during any future reboot.

Step 8 As noted in Defining and Applying Policy Templates, page 8-22, there are a few parameters that can only be configured via the controller web interface or CLI. If any such changes are required, make them at this time via either the controller web interface or a CLI session and save the controller configuration to nonvolatile (flash) memory.

Performing Tasks Across Multiple WLAN Controllers

As has been described in other sections of this document, WCS makes it easy to update WLAN controller resident software such as operating software and IDS signatures from either WCS itself or other TFTP servers in the network. When working with multiple WLAN controllers, especially during a network upgrade or other deployment, WCS makes it possible to perform these types of tasks and others for a group of selected controllers with a minimal amount of keystrokes. The maximum number of WLAN controllers that can be selected for these operations is currently set at the display page size of 20 controller entries. All selected WLAN controllers must be resident on a single WCS display page using the **Configure > Controllers** menu selection. Controllers cannot be included in a selection set if the desired controllers are found on multiple display pages.

For example, if two controllers are being installed at six sites that are located in three different regions of a country, this capability can be used to initiate a download of the latest controller operating software to each of the controllers in each region. Instead of having to initiate twelve separate WCS command sequences to get this done, the ability to specify multiple controllers allows it to be done with three (each sequence specifying a download to all controllers in a single region from a regional TFTP server). If you want to initiate a save to configuration and reboot for all twelve controllers, this can be done via two WCS commands instead of twelve. During a deployment where many controllers may need many of the same functions performed repeatedly, the ability to use WCS to direct such actions at multiple controllers across the management domain saves time and work on the part of the central administration staff.

The types of functions that can be targeted at multiple controllers in this manner include the following:

- Updating operating software to multiple controllers—Allows the administrator to load new controller operating firmware to multiple controllers that have been successfully added to WCS and are currently reachable. It can be accessed via **Configuration > Controllers**, selecting the target controllers from the list presented on the screen and proceeding with the steps detailed in Managing Controller Operating Software, Web Authentication Bundles, and IDS Signatures, page 8-28. It can also be performed on all controllers in a configuration groups via **Configuration > Controllers > Config Groups**.
- Downloading IDS signatures to multiple controllers—Allows the administrator to load new IDS signatures to multiple controllers that have been added to and are currently reachable by WCS. This feature is available by accessing **Configuration > Controllers**, selecting the target controllers from the list presented on the screen, and following the procedure indicated on the screen. It can also be performed on all controllers in a configuration groups via **Configuration > Controllers > Config Groups**.
- Downloading customized web authentication to multiple controllers—Allows the administrator to download a customized web authentication page to multiple controllers. This feature is available by accessing **Configuration > Controllers**, selecting the target controllers from the list presented on the screen and following the procedure indicated on the screen. It can also be performed on all controllers in a configuration groups via **Configuration > Controllers > Config Groups**.
- Saving configuration for flash to multiple controllers—Allows the administrator to initiate writing the current configuration file to nonvolatile (flash) memory on several controllers simultaneously. It is a very useful feature especially immediately after applying policy templates to multiple controllers, because these controllers would normally require this command to be issued individually to save updated configuration information to nonvolatile memory. This feature is available by accessing **Configuration > Controllers**, selecting the target controllers from the list presented on the screen by enabling their checkboxes and then selecting "Save Config to Flash" from the drop-down command selector in the upper right-hand portion of the screen.

  This function can also be performed on all controllers in configuration groups via **Configuration > Controllers > Config Groups**.
- Refreshing WCS configuration from multiple controllers—This feature (described in detail in Non-Selective Synchronization, page 8-36) allows the administrator to initiate the refreshing of the stored configuration contained in the WCS databases from multiple controllers simultaneously. This is useful in correcting situations where it is suspected that the WCS database is no longer in sync with the configuration contained in the controller for multiple categories of configuration objects. Such a situation can result, for example, when changes are made to the WLAN controller configuration in WCS but because of a communication or other error in the controller, these changes did not take effect in the device. A refresh of WCS configuration from multiple controllers can also be performed on all controllers in a configuration groups via **Configuration > Controllers > Config Groups**.
- Initiation of a re-boot in multiple controllers—This feature is also accessible from **Configuration > Controllers** and allows multiple controllers to be re-booted at once. It is a useful feature when updating operating software in multiple controllers, because these updates do not become effective until each controller is re-booted. To initiate a re-boot of multiple controllers, access the **Configuration > Controllers** screen, select the target controllers from the list presented by enabling their check boxes, and then selecting **Reboot Controllers** from the drop-down command selector in the upper right-hand portion of the screen.
- Note that all controllers that are part of a configuration group can be rebooted either sequentially (cascade reboot) or in parallel via the reboot menu tab in **Configuration > Config Groups >*configgroupname*** as shown in Figure 8-18. Rebooting the controllers in a configuration group via this method is an individual action; that is, it does not require that templates be applied to the controllers.

- Configuration Backup Scheduled Task—This feature is useful both as a scheduled task that is executed routinely to archive copies of all known (and currently reachable) controller configurations as well as a method to initiate an immediate archival of those same controller configurations.

Note For more information, see 13.6 CSCsd54800—Config Backup Scheduled Task Fails if One Controller is Unreachable.

Unlike other methods of performing configuration backups, a snapshot of *all* controller configuration in the entire network can be obtained simply by running this one scheduled task. Running this task before making major changes to widespread controller configurations in your network is typically a good idea. See Configuration Backup, page 8-117 for more information about this useful configuration management utility.

Deployment Models

This section discusses two basic deployment models for WCS in the enterprise:

- Campus deployment
- Branch deployment

Note that all illustrations in the following sections have been simplified to focus primarily on the deployment of wireless network management. As such, these illustrations may not illustrate in detail all recommended wired or wireless infrastructure components as specified in other chapters of this SRND.

Campus Deployment

The most common campus deployment model for WCS is as a centralized component managing multiple WLAN controllers (one or more combinations of WLC, WLCM, or WiSMs) interconnected via a modern high-speed campus local area network as shown in Figure 8-72 (optional location appliance is also shown here).

Figure 8-72 Campus WCS Deployment using Single WCS and Location Appliance

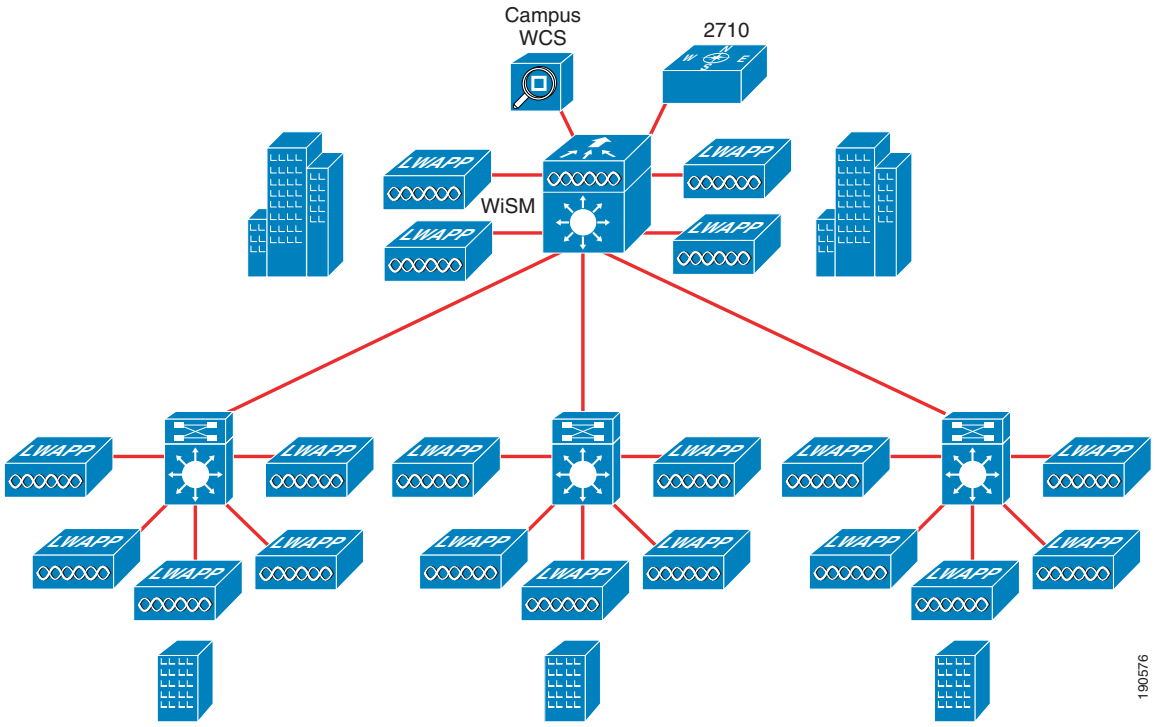

In this model, up to 3000 lightweight access points dispersed among a maximum of 250 WLAN controllers are supported as of WCS version 4.0. (This is dependent on the type of system used for WCS deployment and other factors. See *Cisco Wireless Control System Release Notes, Release 4.0* for further details about WCS capacities.)

This is not a "hard" limitation but rather a limit to which the solution has been tested and is supported by Cisco Systems.

The campus management model illustrated in Figure 8-72 illustrates the use of the Catalyst 6500 Wireless Services Module (WiSM) at the main campus site that provides WLAN controller services for lightweight access points in the main building complex as well as the other locations shown. This use of a Catalyst 6500 WiSM controller module is not mandatory, and the same design model can be applied using external WLAN controllers located centrally instead if so desired.

The approach shown in Figure 8-72 has been found to be suitable for the vast majority of customer wireless campus deployments. In combination with well-designed modern high speed LANs, it provides excellent performance with a highly scalable centralized management base (including optional location-based services) that can be easily scaled without requiring an extensive re-design of the network as the enterprise grows.

In the few cases of very large campuses, this model can be scaled for even greater capacity by grouping WLAN controllers and placing each group under the control of a different WCS server. This approach is known as the *separation of management domains* and is illustrated in Figure 8-73.

Figure 8-73 Multiple WCS/Location Appliance Campus Deployment

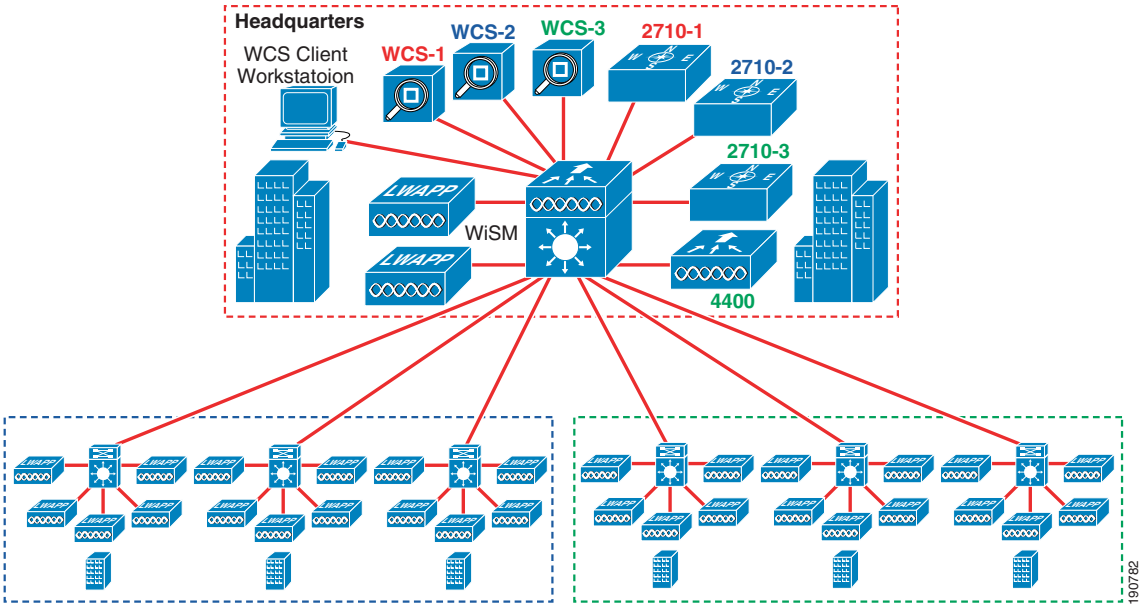

Three WCS servers (WCS1, WCS2, and WCS3) and three location appliances (2710-1, 2710-2, and 2710-3) are co-located at the campus data center and main office building. Each of these WCS/Location Appliance pairs has been configured such that their management and location domains coincide with a different portion of the overall campus (as shown by the red, blue, and green outline boxes). Color-coded text has been used in the illustration provide clarity, indicating that each of the two controllers on the WiSM module are assigned to the management domain of WCS-1 and WCS-2 respectively with the standalone 4400 WLAN controller assigned to WCS-3. Note that each and every WLAN controller is defined to a single WCS management domain (and therefore only one WCS management server).

One or more client workstations (shown in Figure 8-73) can access any of the WCS servers. Each management server can be accessed from a single client workstation using multiple browser windows. It is not uncommon to see client workstations used precisely for this purpose located at the main campus building complex, typically within a network operations center (NOC).

Each WCS server can be configured to generate e-mail notifications (which can be relayed as pager and cell phone text messages) to NOC management staff personnel informing them of critical alarms that have been generated within any of the management domains. NOC personnel can then respond by accessing the proper WCS server from WCS client workstations to investigate and rectify the alarm situation.

To assure that NOC personnel have visibility to all traps issued by WLAN controllers in any of the management domains, the IP address of an overall enterprise network management system such as HP OpenView, Tivoli, and so on, can be entered as an additional trap receiver for each WLAN controller. This assures that all traps can be seen in a central location (in addition to each individual WCS), which facilitates problem resolution in organizations so equipped. Each WLAN controller can also be configured with the address of a NOC remote syslog server. In a similar fashion, this practice provides NOC personnel with visibility to all syslog messages generated by WLAN controllers in each management domain.

As shown in Figure 8-73, the use of a modular approach allows you to scale the campus model beyond the controller and access point limitations mentioned previously. This provides the ability to manage much larger (albeit much less common) campus networks than would otherwise be the case with a single WCS server but still retain a design that is supported by Cisco.

Branch Deployment

This section describes scenarios where the main corporate campus comprises only a minority presence in terms of installed wireless infrastructure within the enterprise. In these cases, the majority of wireless infrastructure as well as mission-critical wireless usage are found in *remote branch offices*.

Figure 8-74 shows a typical WLAN management deployment model in a network servicing remote branch offices (with optional location appliance included). This design illustrates the most common scenario of a single WCS server at a central headquarters location. As can be seen from the information contained in the *Cisco Wireless Control System Release Notes, Release 4.0*, depending on the choice of hardware and network capacity considerations, this WCS server can support up to 3000 lightweight access points distributed over 250 WLAN controllers (keep in mind that the limit on total tracked devices in the location appliance is 2500).

Figure 8-74 Remote Management of Branch Offices—Single Centralized WCS Server

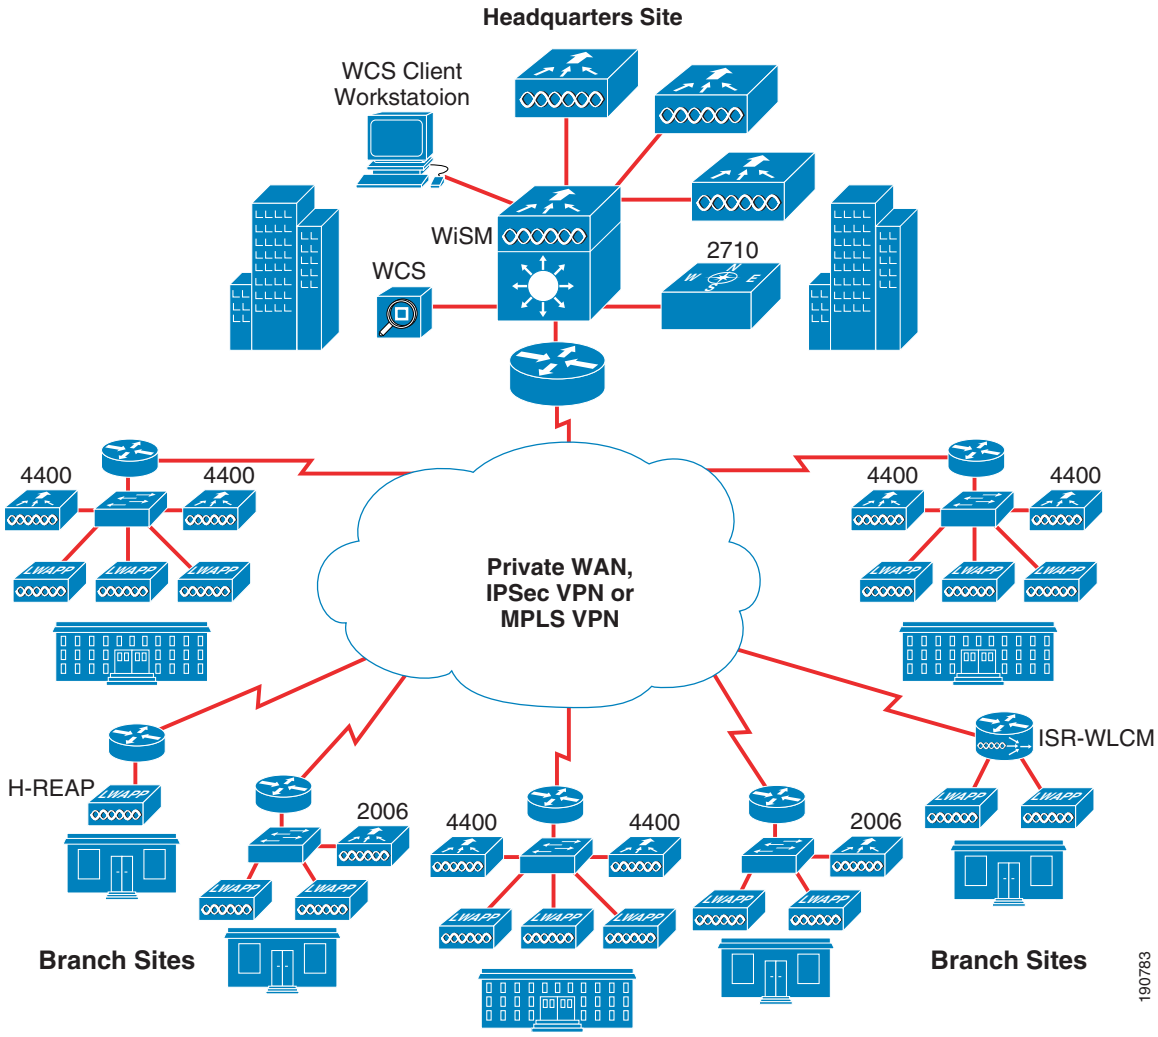

In this type of network design, application computing resources may be located centrally with backup WAN services provisioned in the event of a network outage. In the case of mission-critical application computing resources, these resources may be located in the branch (such as for point-of-sale systems in a retail environment) so as to withstand a complete interruption of primary and backup WAN service. In

the event of a WAN connectivity failure, management and location services are lost but basic wireless connectivity to local resources are preserved. (This availability is for wireless traffic that does not depend on a AAA server located at the central headquarters site, or traffic that is using a AAA server located in the branch.)

Non-mission critical systems (such as accounting systems, personnel records, and so on) are usually centrally located in both cases. As shown in Figure 8-74, the choice of WLAN controller at each site can vary. This can range from a Hybrid REAP (H-REAP) implementation designed to service sites requiring no more than two or three H-REAP lightweight access points per site to those with as many as 50 lightweight access points per site or more, which requires one or more 4400-series WLAN controllers to be deployed.

In some cases, the number of branch locations as well as the number of lightweight access points may be greater than the capacity of a single WCS server, even when deployed on the most robust available hardware. The network illustrated in Figure 8-75 illustrates such a case, where either the total number of lightweight access points is greater than 3000 or the total number of deployed controllers exceeds 250. In those situations, multiple centralized WCS servers can be deployed at the central site, splitting the network into multiple management domains (and multiple location domains).

Figure 8-75 Remote Management of Branch Offices—Multiple Centralized WCS Servers

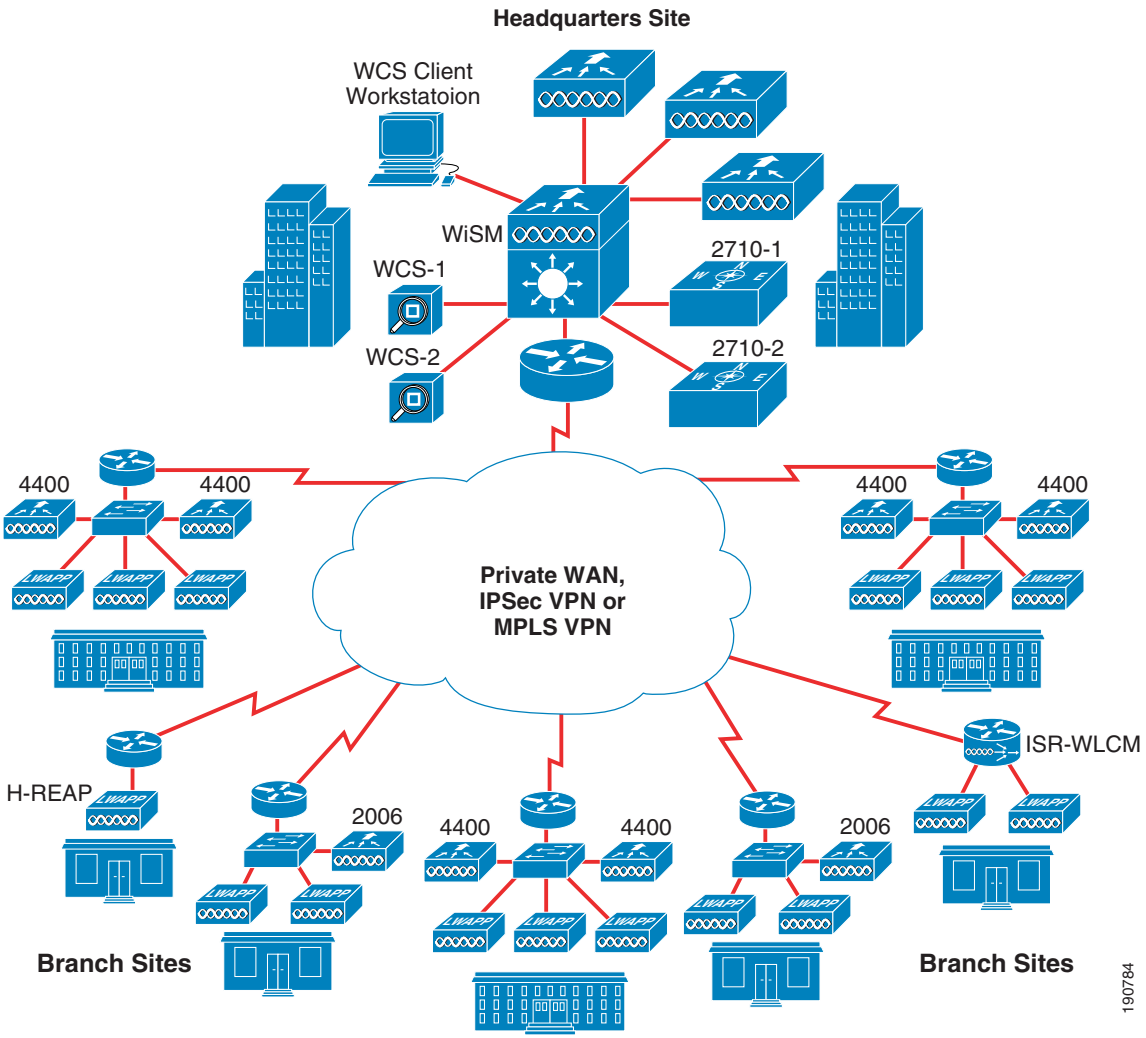

The allocation of WLAN controllers between the WCS servers in Figure 8-75 is the choice of the designer; however, a common approach is to partition the design into at least two WCS management domains. The first management domain you might specify can be for the wireless infrastructure that resides at the headquarters site. One or more additional management domains can then include the branch sites. WCS servers and location appliances can then be allocated to these management domains as appropriate. These WCS servers are then accessible from WCS client workstations located at the headquarters site.

Although the design described in Figure 8-75 may provide satisfactory performance for the majority of large scale wireless branch networks, there are alternatives for even larger (albeit much less common) branch networks that may make more sense depending on the organization itself, the way it is structured, and its underlying WAN topology. Some organizations structured along regional boundaries may find that a network infrastructure more representative of their structure is a better overall fit. In these cases, for example, it is not uncommon to find a central headquarters site along with several regional headquarters sites containing corporate resources in both locations. Figure 8-76 illustrates one such regionalized approach for an enterprise that is sub-divided into four regions plus a headquarters campus site.

Figure 8-76 Remote Management of Branch Offices—Multiple Regionalized WCS Servers

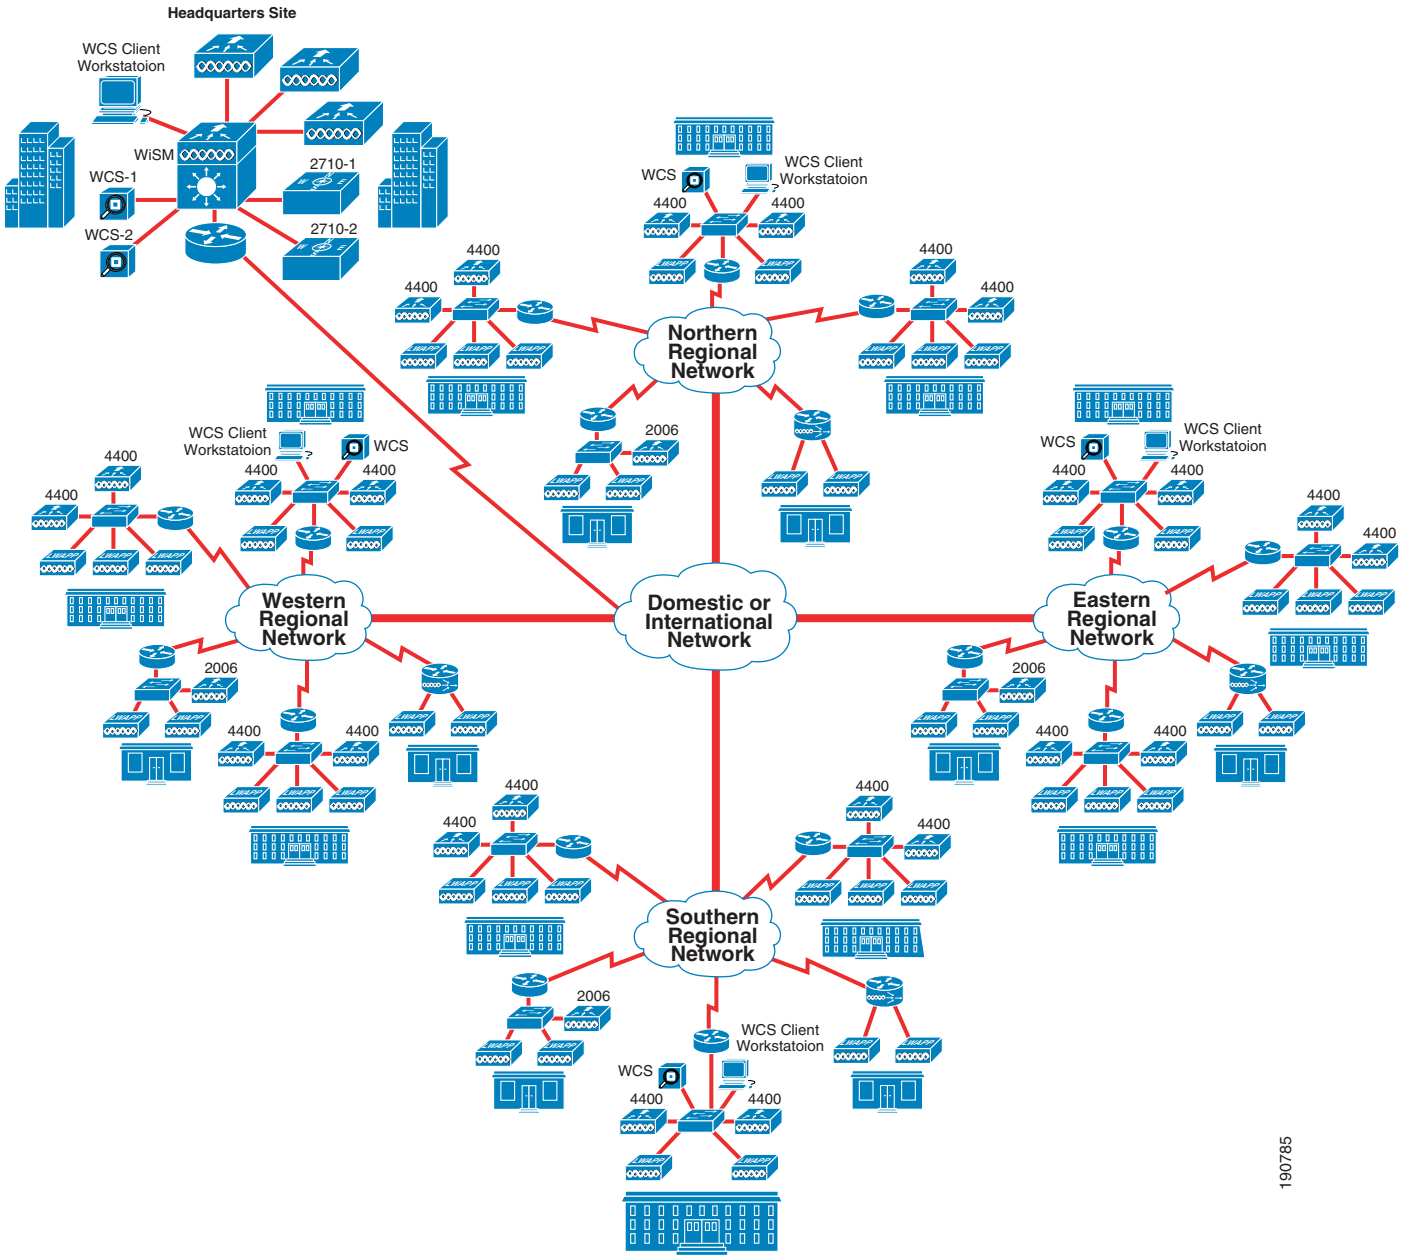

The branch sites in each region are network-managed by regional network operations staff located at each respective regional headquarters location. These regional headquarters locations are also full-time branch locations that, from a business operational standpoint, are peers with the other branch locations in the region. Regional connectivity is provided by one of the WAN networking alternatives already discussed, and the regional networks are connected via very high bandwidth connections to a corporate WAN network that is provided by one or more major service providers.

At the corporate headquarters site, a WCS server and location appliance are present to provide management and location services for the headquarters campus only (these do not interact with branch sites). Headquarters NOC staff can directly manage the wireless network in any branch location via their WCS client workstations (which have access to the WCS server in any region). Each regional WCS

server can be configured to generate e-mail notifications (which in turn can translate into text message notifications to cell phones, pagers, and PDAs) to headquarters network management staff, informing them of critical alarms that have been generated within the region. Headquarters personnel can then respond appropriately via their global WCS client workstations.

The IP address of a headquarters-based enterprise network management system (NMS) can be entered as an additional trap receiver for each WLAN controller in all sites. This ensures that headquarters personnel are made aware via the enterprise NMS of any conditions that trigger trap generation in any of the branch sites.

Note Note that simply defining the IP address of the headquarters WCS server as a trap receiver in the branch WLAN controllers does not provide the desired visibility to branch WLC traps. The headquarters WCS server ignores incoming traps from WLCs that have not been explicitly defined to it via **Configure > Controllers > Add Controller**. Adding a branch controller to the headquarters WCS server then enables polling of the controller by the corporate WCS, which negates one of the advantages of the regional design (reducing polling traffic).

Each branch WLAN controller can be configured with the address of a remote syslog server at the headquarters site to ensure that headquarters personnel have visibility to syslog messages generated by any of the branch WLAN controllers.

The regionalized approach shown in Figure 8-76 and other regionalized designs like it can offer the following key advantages to those customers whose organizational makeup and size allows them to make good use of it:

- A network failure in any one of the regional networks or in the corporate network itself is unlikely to cause a loss of WCS management in the branches outside of the affected region. From a network management perspective, a corporate network interruption affects only the ability of the corporate NOC staff to receive trap, syslog, and e-mail updates as well as their ability to directly manage resources in any of the branch sites. Depending on the degree to which resources are regionalized, each region can retain a fairly high degree of operational autonomy in spite of a corporate network disruption.
- Polling between WCS and each WLAN controller (and optionally between the location appliance and each controller) is confined entirely to each region and does not occur across the corporate network, which may be of interest to those planning to deploy large wireless networks. Assuming an even distribution of polling traffic among regions, the total traffic volume in any of the regional networks would be approximately estimated at only about 25 percent of what would be seen across a single network with all management servers centralized at one location.

Thus far, two remote branch management deployment models have been discussed that provide a very workable solution for the great majority of all wireless network management needs. Even so, there are still some organizations whose sheer scope and size may make even a regionalized solution such as shown in Figure 8-76 less than optimal. For this small group of extraordinarily large (in many cases multi-national) entities, you can institute WCS management at lower levels of the network, further down below even the regional level.

Some of the specialized concerns in these extremely large networks that can require such designs include the following:

- Organizational preferences—Some organizations may grant individual branch offices greater operational autonomy than might otherwise be the case. Although data and programs for many administrative background processes may be contained on corporate or regional servers, computing and network resources deemed “mission-critical” may be located at each branch site. This allows

the branch to function with not only a high degree of managerial empowerment but with the capability to act almost as a standalone autonomous business unit when severe environmental events preclude any form of WAN connectivity to corporate or regional headquarters.

In this type of organization, it may be preferred for the local branch to have the capability of managing its wireless LAN infrastructure in an autonomous fashion more analogous to its operational capabilities, even when situations make it impossible to establish external network connectivity. In both the centralized as well as the regional approaches illustrated in Figure 8-74, Figure 8-75, and Figure 8-76, this type of outage requires local branch personnel lacking access to their WCS servers to manage their WLAN controllers directly via the web or command line interfaces. No location services capabilities are available during any such outage because of the lack of accessibility to WCS and the location appliances located at the corporate or regional headquarters (this includes on-demand location).

- Mission-critical location-based services—Some very large branch enterprises may make use of location-based services applications at each branch location such that normal business functions may be impacted severely if LBS is not available. Looking back at the designs in Figure 8-74, Figure 8-75, and Figure 8-76, the location appliance is located either at a central or regional headquarters site. Any third-party location client-server applications that are located within the branch are not able to display current location information if the location appliance becomes unavailable because of a prolonged WAN failure. An example is a large national hospital corporation that depends on its LBS system to quickly locate critical medical equipment. A sudden disruption in such a system may impact the level of service that the hospital is able to deliver to its patients.
- Polling traffic—Taking into the consideration the information presented in Device Status Polling, page 8-107, there is a potential for very large enterprise networks to produce a significant amount of network polling traffic, especially at the central headquarters or regional headquarters where WCS servers happen to be located. Depending on the current use of existing circuits, this added traffic between WCS and WLAN controllers (as well as between the location appliance and WLAN controllers) can be of concern.

In the few extremely large-scale deployments where these areas must be addressed, the added cost of deploying a WCS server (and location appliance) within each branch may be justified. This approach eliminates the impact of management and location services polling upon the WAN completely by relocating the WCS servers and location appliances to each individual branch. It also adds full management and location-based services survivability to the branch in spite of prolonged WAN interruption, thereby allowing branch wireless management autonomy. In the event of a WAN interruption at either the regional or corporate level, local branch WLAN users continue to have access to branch-resident mission-critical resources located, and local branch technical personnel have complete and unfettered management access to their wireless LAN infrastructure as well.

Traffic Considerations When Using WCS in Large Networks

For customers wishing to deploy large or very large networks, an understanding of the traffic volumes that are involved when routine management polling occurs can be useful in making proper design choices.

Traffic Sources

In a system consisting of a WCS, location appliance, WLAN controllers, and lightweight access points, the following categories are the main sources of network traffic:

- Between WLAN Controllers and WCS:

  - Device status polling
  - Client statistics polling
  - RF statistics polling
  - Rogue access point polling
  - Configuration audit reports and network audits
  - Controller configuration refresh (including cold-start refresh)
  - WCS configuration refresh
  - Controller configuration backups
  - Software, configuration and IDS signature downloads
- Between WLAN controllers and the location appliance:
  - Client polling
  - Asset tag polling
  - Rogue AP/client polling
  - Statistics polling
- Between WCS and the location appliance:
  - Network design synchronization
  - Location appliance backup

WLAN Controllers and WCS

WCS obtains information about the status of WLAN controllers, lightweight access points, WLAN clients, asset tags, and rogues in two ways: via WCS-initiated SNMP polling of the WLAN controllers, and unsolicited SNMP traps generated by WLAN controllers. As mentioned in previous sections, WCS does not poll lightweight access points, WLAN clients, asset tags, or rogue devices directly, but instead relies on information that WLAN controllers proxy to WCS about these entities.

In WCS, configure the various categories of polling as well as the regularity of polling via one of four scheduled tasks:

- Device status polling
- Client statistics polling
- Statistics polling
- Rogue AP polling

The following sections discuss each of these in more detail and provide a brief understanding of the level of traffic impact each polling categories can assess on total network management traffic.

Although SNMP polling makes up the predominant portion of the total management traffic between WCS and WLAN controllers, it is by no means the exclusive source of management traffic. SNMP traps flowing from controllers to WCS and other trap receivers also contribute a traffic component, albeit one that is rarely a concern in most cases. The number of traps enabled, the number of WLAN controllers, the amount of times additional SNMP polling is triggered by the reception of traps, the number of trap receivers, and the frequency at which trap-generating events occur all dictate the impact of SNMP traps on overall management traffic.

Utility functions such as backup and restores of controller and WCS database configurations can also contribute to network congestion, primarily in very large networks when such functions might be performed on many controllers simultaneously. Taken individually, such actions are typically of minor consequence, but when initiated on large groups of controllers simultaneously, it would not be unusual to see traffic spikes that may become noticeable overall. An example of where this may be commonly seen is with the case of controller software downloads (see Managing Controller Operating Software, Web Authentication Bundles, and IDS Signatures, page 8-28) and controller configuration backups that tend to be the most traffic intensive of the utility functions. Traffic because of other utility functions such as configuration refreshes to and from WLAN controllers (described in Non-Selective Synchronization, page 8-36) are usually not significant and typically should not pose a problem.

Reporting functions (such as the configuration audit reports) cause WCS to inventory the configuration of each selected controller and compare the results against the contents of the WCS database. Configuration Audit Reports and Network Audits, page 8-113 discusses the best practices surrounding the running of these reports and also some detailed analysis of the traffic flows that are involved for both SNMPv2c and SNMPv3. The network audit scheduled task (Network Audit, page 8-118) is a derivative of the configuration audit report except that it runs against all controllers defined to WCS instead of a subset of selected controllers. In large networks comprising many WLAN controllers and lightweight access points, it is always good practice to consider running the network-wide network audit scheduled task at times of low network usage.

Using the discussions of initial and incremental traffic as a guide (described in Device Status Polling, page 8-107 through SNMP Traps, page 8-111), it is possible to gain a general understanding of the traffic volume that might be generated by the various management polling methods in a proposed design. Note that estimates stated in the subsections to follow are based on lab testing under closely controlled circumstances. The number of devices, clients, and rogues detected in your network will likely vary between polling cycles.

This information is provided so that designers and architects of network systems can use the traffic volume information in conjunction with information they already know about network utilization in their own environment to make intelligent design decisions. It is not provided with the intent of serving as a precision prediction tool, but rather to educate readers with large networks as to the magnitude of potential management traffic volumes they may incur in an effort to enable better overall network designs.

Management traffic volumes should be considered not only in conjunction with the present network utilization, but with those traffic loads that can be reasonably expected in the near future because of growth and expansion. If location appliances are to be included in the design, a separate analysis should be conducted focusing on polling traffic between the location appliance and the WLAN controllers (see WLAN Controllers and the Location Appliance, page 8-115).

In the minority of large wireless networks where the impact of SNMP polling on network utilization is deemed to be excessive, the polling intervals for device status, client statistics, radio statistics, and rogue access point polling can be adjusted as described in *Cisco Wireless Control System Release Notes, Release 4.0*, which is available at the following URL: http://www.cisco.com/en/US/products/ps6305/prod_release_note09186a00806b0811.html.

By staggering and increasing the polling intervals, you can better distribute the aggregate volume of polling traffic that is being introduced into the network and lower increases in peak network utilization that are because of management SNMP polling.

Unless otherwise noted, all traffic quantities in the following sub-sections are UDP byte and packet counts for a single controller that were measured on a 10/100/1000 Ethernet LAN using Ethereal 0.99.0. Measurements were taken for both bidirectional and unidirectional traffic flows, along with a measurement of the total elapsed time required for the polling process to complete. Hardware used in these tests were Cisco Catalyst c3750 Ethernet switches and 4400-12 WLAN Controllers with 4.0.155.5 controller software and AP1242 access points.

Device Status Polling

Device status polling is conducted in WCS by the Device Status scheduled task. WCS is configured by default to conduct device status polling every five minutes. In addition to being responsible for updating device reachability, device status polling provides WCS with information such as the following:

- Controller SysUpTime, total memory, free memory, CPU utilization, and operating software version
- Lightweight access point and radio interface administrative status
- Coverage, load, and interference profile status
- 802.11 privacy options in use
- Beacon periods currently in use
- MAC addresses of AP neighbors

The amount of traffic exchanged during device status polling is not affected by the presence of wireless LAN clients, asset tags, rogue access points, or rogue clients. The volume of polling traffic produced was seen to be tied to the number of lightweight access points registered to the WLAN controller.

Figure 8-77 provides average SNMP v2c/SNMPv3 UDP traffic volumes observed during testing between WCS and a 4400-12 WLAN controller. The tables provide us with baseline device status polling traffic figures for a single 4400-12 controller with zero access points registered, and allows visualization of how this traffic flow increases with additional registered access points.

The table in Figure 8-77 is broken down into ranges. Note that within each range, the traffic quantities listed represent the amount of traffic observed when the number of registered access points was within the range. For example, when device status polling occurs using SNMPv2c between WCS and a controller with four registered infrastructure access points, on average an exchange of approximately 4805 bytes of UDP data was observed each time WCS polled the WLAN controller for device status information. However, if the number of registered infrastructure access points increases to nine, the average number of bytes exchanged between WCS and the WLAN controller increases to approximately 7154 during each iteration of device status polling.

Although not intended as a precision prediction tool, this information can be useful in understanding how device status polling traffic might grow beyond the twelve access points shown.

Figure 8-77 Device Status Polling Traffic

SNMP v2c	0 AP	1 - 4 AP	5 - 9 AP	10 - 12 AP
WCS↔WLC, average bytes	4275	4805	7154	10443
WCS↔WLC, average seconds	0.051	0.049	0.055	0.065
WCS↔WLC, average packets	26	26	30	36
WCS→WLC, average packets	13	13	15	18
WCS→WLC, average bytes	1226	1226	1549	1990
WLC→WCS, average packets	13	13	15	18
WLC→WCS, average bytes	3049	3579	5605	8453

SNMPv3	0 AP	1 - 4 AP	5 - 9 AP	10 - 12 AP
WCS↔WLC, average bytes	6967	7432	10288	14051
WCS↔WLC, average seconds	0.09	0.084	0.096	0.138
WCS↔WLC, average packets	28	28	32	38
WCS→WLC, average packets	14	14	16	19
WCS→WLC, average bytes	2552	2509	3060	3777
WLC→WCS, average packets	14	14	16	19
WLC→WCS, average bytes	4415	4923	7227	10274

190786

Client Statistics Polling

Client statistics polling is conducted in WCS by the clients statistics polling scheduled task (Client Stats Poll). WCS is configured by default to conduct client statistics polling every fifteen minutes. Some of the information that is gathered by WCS via client statistics polling includes but is not limited to the following:

- Average SNR and RSSI of clients
- Number of packets received and sent from/to the client
- Number of bytes received and sent from/to the client
- Number of policy errors that have occurred for the client
- Client status

The amount of traffic exchanged during client status polling is not related to the number of lightweight access points registered to the WLAN controller. It is also not affected by the presence of Layer 2 asset tags, rogue access points, or rogue clients.

Note Asset tags that associate to lightweight access points as WLAN clients (such as PanGo Locator LAN tags) contribute to the amount of traffic produced during client statistics polling. AeroScout asset tags do not associate to access points.

The traffic volume is driven primarily by the number of WLAN clients that have associated to the lightweight access points serviced by the controller being observed. Lightweight access points that do not have WLAN clients associated do not contribute to client statistics polling traffic.

The distribution of clients between lightweight infrastructure access points was seen to have little if any impact on the aggregate amount of client statistics polling traffic produced for that controller. For example, four WLAN clients associated to a single lightweight access point affiliated with a controller result in approximately the same volume of client statistics polling traffic between that controller and WCS as four WLAN clients, each associated to individual lightweight access points on the same controller.

Figure 8-78 provides average SNMP v2c/SNMPv3 UDP traffic volumes observed during testing between WCS and a 4400-12 WLAN controller with a single registered lightweight access point. The tables provide baseline client statistics polling traffic figures for a single 4400-12 controller with zero clients associated, and allows visualization of how this flow might increase with additional clients.

Figure 8-78 Client Statistics Polling Traffic

SNMP v2c	0 clients	1 - 3 clients	4 - 6 clients
WCS↔WLC, average bytes	3950	3950	6118
WCS↔WLC, average seconds	0.028	0.012	0.033
WCS↔WLC, average packets	8	8	10
WCS→WLC, average packets	4	4	5
WCS→WLC, average bytes	892	892	1618
WLC→WCS, average packets	4	4	5
WLC→WCS, average bytes	3058	3058	4500

SNMPv3	0 clients	1 - 3 clients	4 - 6 clients
WCS↔WLC, average bytes	4502	4510	6776
WCS↔WLC, average seconds	0.023	0.022	0.048
WCS↔WLC, average packets	8	8	10
WCS→WLC, average packets	4	4	5
WCS→WLC, average bytes	1258	1264	2088
WLC→WCS, average packets	4	4	5
WLC→WCS, average bytes	3244	3246	4688

190787

The table in Figure 8-78 is divided into ranges. Note that within each range, the traffic quantities listed represent the amount of traffic observed when the number of clients associated to registered access points is within the range. For example, when client statistics polling occurs using SNMPv2c between WCS and a controller with three associated clients, on average approximately 3950 bytes of UDP data are exchanged between WCS and the WLAN controller. However, if the number of associated clients increases to six, the average number of bytes exchanged between WCS and the WLAN controller increases to approximately 6118 during each iteration of client statistics polling.

Although not intended as a precision prediction tool, this information can be useful in understanding how client statistics polling traffic might grow beyond the six clients shown.

Statistics Polling

Statistics polling is conducted in WCS using the Statistics scheduled task. The statistics that are gathered by this scheduled task concern the lightweight access point radio interfaces. WCS is configured by default to conduct statistics polling every fifteen minutes. Some of the information that is gathered by WCS via statistics polling includes but is not limited to the following:

- Radio interface transmit power level
- Radio interface operational status
- Number of WLAN clients associated to a radio interface
- Percentage of time interface radio receiver/transmitter is receiving/transmitting packets
- Channel utilization
- Number of clients with below-threshold SNR
- Status of whether load, coverage, noise, or interference thresholds have been exceeded
- Transmitted and received fragment counts
- FCS error count

The amount of traffic exchanged during statistics polling is not affected by the presence of wireless LAN clients, asset tags, rogue access points, or rogue clients. Rather, the volume of traffic produced is primarily driven by the number of lightweight access points (and the number of 802.11 radios they contain) that registered with the WLAN controller.

Figure 8-79 provides average SNMP v2c/SNMPv3 UDP traffic volumes observed during testing between WCS and a 4400-12 WLAN controller. The tables provide baseline statistics polling traffic figures for a single 4400-12 controller with zero registered lightweight access points and allows you to visualize how this flow of traffic might increase as the number of registered infrastructure dual-band lightweight access points is increased.

Figure 8-79 Statistics Polling Traffic

SNMP v2c	0 AP	1-2 AP	3-4 AP	5-6 AP	7-8 AP	9-10 AP	11-12 AP
WCS↔WLC, average bytes	4225	7259	10279	13229	16179	19135	22036
WCS↔WLC, average seconds	0.014	0.062	0.086	0.157	0.163	0.199	0.269
WCS↔WLC, average packets	12	16	18	20	22	24	26
WCS→WLC, average packets	6	8	9	10	11	12	13
WCS→WLC, average bytes	932	2453	4015	5526	7036	8550	10022
WLC→WCS, average packets	6	8	9	10	11	12	13
WLC→WCS, average bytes	3293	4806	6264	7704	9143	10585	12014

SNMPv3	0 AP	1-2 AP	3-4 AP	5-6 AP	7-8 AP	9-10 AP	11-12 AP
WCS↔WLC, average bytes	5459	8737	11722	14678	17634	20590	23517
WCS↔WLC, average seconds	0.060	0.059	0.097	0.137	0.173	0.230	0.328
WCS↔WLC, average packets	14	18	20	22	24	26	28
WCS→WLC, average packets	7	9	10	11	12	13	14
WCS→WLC, average bytes	1637	3293	4836	6350	7864	9378	10863
WLC→WCS, average packets	7	9	10	11	12	13	14
WLC→WCS, average bytes	3822	5444	6886	8328	9770	11212	12654

190788

The table in Figure 8-79 is divided into ranges. Note that within each range, the traffic quantities listed represent the amount of traffic observed when the number of registered dual-band access points was within the range. For example, when statistics polling occurs using SNMPv2c between WCS and a controller with four registered dual-band access points, on average approximately 10,279 bytes of UDP data are exchanged each time WCS polls the WLAN controller. However, if the number of registered dual-band access points increases to six, the average number of bytes exchanged between WCS and the WLAN controller increases to approximately 13,229 during each iteration of statistics polling.

Although not a precision prediction tool, this information can be useful in understanding how polling traffic might grow beyond the 12 access points shown.

Rogue Access Point Polling

Rogue access point polling is conducted in WCS via the Rogue AP scheduled task. WCS is configured by default to start rogue AP polling of all controllers every 120 minutes. Some of the information that is gathered by WCS during rogue access point polling includes but is not limited to the following:

- Rogue AP type
- Rogue AP channel, SNR, RSSI, WEP mode, WPA mode, preamble
- Rogue AP SSID, radio type
- Time stamp of rogue AP initial detection
- Total number of rogue clients
- Rogue AP on network status
- Rogue AP containment level
- Total number of detecting APs

- Detecting AP names and MAC addresses

The amount of traffic exchanged between WCS and the WLAN controller during rogue AP polling is not affected by the presence of wireless LAN clients or asset tags. It is driven primarily by the number of rogue access points detected by the lightweight infrastructure access points.

Figure 8-80 provides average SNMP v2c/SNMPv3 IPv4 traffic volumes observed during testing between WCS and a 4400-12 WLAN controller. The tables provide baseline rogue polling traffic figures for a single 4400-12 controller with zero detected rogue APs through 29 detected rogue access points, and allows you to visualize how this flow of traffic might increase as the number of detected rogue access points increases.

Figure 8-80 Rogue AP Polling Traffic

SNMP v2c	0 Rogue APs	1 - 9 Rogue APs	10 - 19 Rogue APs	20 - 29 Rogue APs
WCS↔WLC, average bytes	2592	2592	5158	6986
WCS↔WLC, average seconds	0.013	0.011	0.020	0.019
WCS↔WLC, average packets	5	5	8	11
WCS→WLC, average packets	2	2	3	4
WCS→WLC, average bytes	335	337	928	983
WLC→WCS, average packets	3	3	5	7
WLC→WCS, average bytes	2257	2255	4230	6003

SNMPv3	0 Rogue APs	1 - 9 Rogue APs	10 - 19 Rogue APs	20 - 29 Rogue APs
WCS↔WLC, average bytes	3885	3885	6221	8027
WCS↔WLC, average seconds	0.162	0.180	0.171	0.177
WCS↔WLC, average packets	11	11	14	14
WCS→WLC, average packets	5	5	6	6
WCS→WLC, average bytes	972	972	1387	1530
WLC→WCS, average packets	6	6	8	8
WLC→WCS, average bytes	2913	2913	4834	6497

190789

The table in Figure 8-80 is divided into ranges of detected rogue access points. Note that within each range, the traffic quantities listed represents the amount of traffic that was observed when the number of detected rogue access points was within the range. For example, when rogue AP polling occurs using SNMPv2c between WCS and a controller with nine detected rogue access points, on average approximately 2592 bytes of IPv4 data was exchanged between WCS and the WLAN controller each time WCS polls the WLAN controller. However, if the number of detected rogue access points rises to 12, then the average number of bytes exchanged increases to approximately 5158 on each polling iteration.

Although not intended as a precision prediction tool, this information can be useful in understanding how rogue polling traffic might grow beyond the 29 detected rogue access points tested here.

Note that the presence of associated rogue clients was not found to add any appreciable amount of traffic to the results recorded in Figure 8-80.

SNMP Traps

As described in Relationship Between Traps, Events, Alarms, and Notifications, page 8-69, a trap is a notification issued by a managed device to the network management station (WCS) when a significant event occurs at the managed device. WLAN controllers can be configured to send SNMP traps to up to six trap receivers. In contrast to a response to a polling request, the information contained in a trap is typically sent in an unsolicited manner. When traps are configured and enabled in the managed device (WLAN controller), they are sent to WCS as the events that generate them occur.

Traps are enabled or disabled via **Configure > Controllers > Management > Trap Control**, and trap receivers are defined using **Configure > Controllers > Management > Trap Receivers**, as described in Relationship Between Traps, Events, Alarms, and Notifications, page 8-69. Both trap receiver as well as trap control configuration objects can be defined via policy templates. For proper WCS event notification, it is always required that WCS be defined as one of the trap receivers in each WLAN controller comprising the management domain. You may opt to define other trap receivers as well depending on your organizational policy and any other enterprise management systems in use, but always make sure to include a WCS server as a trap receiver for your WLAN controllers.

Because traps are sent without acknowledgement, it is possible that a transmitted trap is never received by WCS. This can happen, for example, in highly utilized or congested WAN networks because of packet discard algorithms that take place in the service providers network, or simply because of QoS mechanisms doing their job and discarding low-priority traffic. Because trap delivery is not guaranteed, WCS also polls for some of the same information that is available via the trap mechanism. Examples include rogue AP traps and various traps pertaining to the state of the lightweight access point radios. The WLAN controller makes much of this same information available to WCS via rogue access point and statistics polling, thereby ensuring that if this information is not received via the trap mechanism, WCS still learns of any extraordinary conditions during the next polling cycle.

The traps sent from WLAN controllers to WCS can vary in size (during lab testing traps were analyzed that ranged in size from approximately 120 bytes to almost 400 bytes). In most cases, the amount of overall traffic added to the network specifically attributable to SNMP traps is inconsequential. However, for those customers that may be considering deploying very large networks with large numbers of controllers and several WCS servers, a few key points should be kept in mind:

- Depending on the specific trap received, WCS may immediately poll the responsible WLAN controller for additional information.

  An example of this can be seen in the AP Registered trap. After reception of this trap, WCS immediately issues a series of poll requests to the controller that initiated the trap for additional device status information. Other examples include traps indicating that the channel country set or power levels have changed. This is a well-known and perfectly acceptable method of obtaining additional information that has not been included in the trap itself, and it allows WCS to gain a better understanding of precisely what the condition is at the WLAN controller that initiated the trap. The point to keep in mind here is that the net contribution to management traffic in this case is more than simply the traffic volume incurred transmitting the trap itself.

- Out-of-date trap receiver lists can increase the level of network trap traffic by causing unnecessary copies of traps to be sent to non-existent stations, stations that are simply not listening on the trap port any longer, or stations that are listening but really should not be provided with the information any longer.

  The amount of trap traffic generated by a controller increases in direct proportion to the total number of trap receivers specified, up to the maximum of six trap receivers (that is, a fully configured controller with six trap receivers defined generates six times more trap traffic that a controller that has only a single WCS configured as a trap receiver and nothing more). In large networks, trap receivers should be specified strictly on a need-to-know basis and only for stations that are actually active trap receivers. Trap receivers used during troubleshooting, testing, or network diagnostics should be promptly disabled or removed promptly after their usefulness has expired and they are no longer necessary. The use of policy templates can make the assignment and prompt removal of any unnecessary trap receivers much easier.

Additional information about the traps configurable via the **Configure > Controllers > Management > Trap Control** page can be found in online help system available under the WCS main menu bar as **Help > Online Help**.

Appendix D, "Examples of SNMP Traps," contains several examples of SNMP traps captured during lab testing. Although not a complete listing of all traps available from Cisco WLAN controllers, this appendix does show actual received traps and decodes much of their content for easier viewing. Complete trap definitions for 4400, 4100, and 2000 series WLAN controllers are defined in MIB files that are available to registered users on the Cisco Connection Online (CCO): http://www.cisco.com.

Configuration Audit Reports and Network Audits

The routine examination of periodic configuration audit reports is a useful tool for the WCS administrator toward ensuring that the integrity of the WCS databases is being verified and maintained. This is especially important in organizations where disparate groups may be responsible for various facets of network operation and maintenance. As is often the case when there are service impacting outages in remote field locations, modifications to WLAN controller configurations may sometimes occur outside the auspices of the network operations center and WCS (in the name of expeditious problem resolution and service restoration).

In these cases, an out-of-sync condition can exist between the configuration stored within the WCS databases and the current configuration of the actual WLAN controller or lightweight access point. Adherence to a policy of routine configuration audit report examination can give advanced warning that such activities have occurred, and more importantly prompt the WCS administrator to the fact that re-synchronization or the re-application of uniform policy templates may be justified.

Configuration audit reports can be run for either a single controller in the management domain (see On-Demand Configuration Audit Reporting, page 8-33), or for all controllers in the management domain via the Network Audit scheduled task (see Scheduled-Task Network Audit Reporting, page 8-34). Although there is typically little concern regarding running the standalone configuration audit report for a single controller, customers with very large wireless networks may wish to consider the traffic impact of the Network Audit task before use during peak periods of network use.

When a configuration audit report is run for a controller, WCS basically retrieves the content of the WLAN controller configuration via SNMP. This is done both for the WLAN controller itself as well as the configuration of any currently registered lightweight access points. The amount of data sent is therefore partially determined by the number of lightweight access points currently registered with the WLAN controller. A small excerpt of the entire SNMP v2c exchange that occurs between a WCS and a 4400-series WLAN controller during a configuration audit can be found in Appendix A, "Excerpt of Configuration Audit Exchange, WCS <-> 4400 WLAN Controller.".

Figure 8-81 provides us with an analysis of the traffic flow observed between WCS and a 4400-12 WLAN controller during the execution of the single controller configuration audit. The number of registered access points varied from zero to the maximum capacity of the controller. Keep in mind that the amount of information transferred during the network audit did not depend on the number of WLAN clients or asset tags associated to the access points registered to the controller, but was very dependent on the complexity of the controller configuration. The data in Figure 8-81 was based on a very minimally (default) configured 4400-12 WLAN controller. More complex controller configurations increases the amount of data transferred during the configuration audit.

Figure 8-81 Configuration Audit Traffic Analysis

SNMP v2c	0 AP	1 AP	2 AP	3 AP	4 AP	5AP	6AP	7AP	8AP	9AP	10AP	11AP	12AP
WCS↔WLC, average bytes	85519	86492	88793	93321	93244	98567	103116	103066	105123	109917	112902	114998	119606
WCS↔WLC, average seconds	1.31	1.73	1.38	1.43	1.69	1.41	1.59	1.43	1.48	1.59	1.63	1.648	1.67
WCS↔WLC, average packets	304	304	306	310	310	316	320	320	322	326	330	332	336
WCS→WLC, average packets	152	152	153	155	155	158	160	160	161	163	165	166	168
WCS→WLC, average bytes	25266	25377	26188	27938	27943	29133	30869	30847	31498	33172	33733	34397	36088
WLC→WCS, average packets	152	152	153	155	155	158	160	160	161	163	165	166	168
WLC→WCS, average bytes	60253	61115	62605	65383	65301	69434	72247	72219	73625	76745	79169	80601	83518

SNMPv3	0 AP	1 AP	2 AP	3 AP	4 AP	5AP	6AP	7AP	8AP	9AP	10AP	11AP	12AP
WCS↔WLC, average bytes	111337	111785	114365	119063	119047	124754	129500	129444	131589	136295	139894	142021	146794
WCS↔WLC, average seconds	1.93	1.45	1.53	1.55	1.64	1.7	1.68	2.28	2.28	2.00	2.10	2.15	2.29
WCS↔WLC, average packets	304	304	306	310	310	316	320	320	322	326	330	332	336
WCS→WLC, average packets	152	152	153	155	155	158	160	160	161	163	165	166	168
WCS→WLC, average bytes	38749	38749	39652	41602	41602	43069	44995	44971	45706	47568	48374	49075	50953
WLC→WCS, average packets	152	152	153	155	155	158	160	160	161	163	165	166	168
WLC→WCS, average bytes	72588	73036	74713	77461	77445	81685	84505	84473	85883	88727	91520	92946	95841

190790

As mentioned previously, the primary use for this information is in gauging the amount of traffic to expect if one is planning on running the all-controller-inclusive network audit scheduled task in very large network configurations. In most such cases, the Network Audit configuration should be scheduled to run during off-peak times of operation or other times when any such audit of all controllers in the management domain would have minimal impact on the users of the network. In those very large networks where there are several WCS servers, each managing different portions of the network over a shared network infrastructure, you may wish to stagger scheduling of Network Audit tasks between the WCS servers such that all WCS servers are not attempting to audit their management domains simultaneously.

Configuration Backup

Lab testing has shown that under version 4.0, a controller configuration archive for a controller with a nearly default configuration is approximately 740,000 bytes in size. When the WLAN controller configuration was archived (either on-demand or via the **Configuration Backup** scheduled task) using SNMPv2c the traffic flows shown in Figure 8-82 were observed over a time period of 13.45 seconds.

Figure 8-82 Single WLC Configuration Archival Traffic Analysis, SNMPv2c

					Packets	Bytes	Packets	Bytes	Avg Packet Size
Address A	Address B	Protocol	Packets	Bytes	A => B	A => B	B => A	B => A	A <=> B
WLAN Controller	TFTP Server	TFTP	2959	893134	1479	825054	1480	68080	302
WCS	WLAN Controller	SNMP	36	4507	18	2180	18	2327	125

190791

The use of SNMPv3 in (Figure 8-83) was shown to have only minimal impact in this case because the bulk of the traffic was because of TFTP and not SNMP components. Elapsed time for the SNMP traffic flow was 13.29 seconds.

Figure 8-83 Single WLC Configuration Archival Traffic Analysis, SNMPv3

					Packets	Bytes	Packets	Bytes	Avg Packet Size
Address A	Address B	Protocol	Packets	Bytes	A => B	A => B	B => A	B => A	A <=> B
WLAN Controller	TFTP Server	TFTP	2959	893134	1479	825054	1480	68080	302
WCS	WLAN Controller	SNMP	36	7711	18	3759	18	3952	214

190792

Designers of very large wireless networks may wish to consider the potential traffic associated with the Configuration Backup scheduled task in their environment because it initiates configuration backups of *all* reachable controllers defined to WCS in a serial fashion.

Note For more information, see section 13.6 CSCsd54800—Config Backup Scheduled Task Fails if One Controller is Unreachable.

Although the peak traffic impact of this task is not high (WLCs are archived one at a time in sequence), it can run for some time depending on the number of the controllers that are reachable. It is good practice to schedule this archival tool for execution during off-peak periods of low usage instead of during peak traffic periods.

Non-Selective Configuration Refresh

Lab testing has shown that the traffic volumes experienced under version 4.0 for the Refresh Config from Controller and the Restore Config to Controller operations are nominal. For SNMPv2c, the traffic flow was observed as 372 packets of 129,000 bytes in a time period of about 1.06 seconds. This traffic was allocated as 186 packets in each direction, with 39,000 bytes from WCS to the WLAN controller and 91,000 bytes returned from the WLAN controller to WCS, with an average UDP packet size of 400 bytes. For SNMPv3, the traffic flow increases to about 372 packets of 160,000 bytes in a time period of about 1.77 seconds. This SNMPv3 traffic was observed to be 186 packets in each direction, 55,000 bytes from WCS to the WLAN controller, and 104,000 bytes returned from the WLAN controller to WCS with an average UDP packet size of 475 bytes.

For Restore Config to Controller, lab testing has indicated that under SNMPv2c, the traffic flow was observed as 1244 packets of 316,000 bytes in a time period of about 6.3 seconds. This traffic was allocated as 622 packets in each direction, with 128,000 bytes from WCS to the WLAN controller and 188,000 bytes returned from the WLAN controller to WCS, with an average UDP packet size of 275 bytes. For SNMPv3, the traffic flow increases to 1244 packets of 427,000 bytes in a time period of about 7 seconds. This SNMPv3 traffic was observed to be 622 bytes in each direction, 184,000 bytes from WCS to the WLAN controller and 243,000 bytes returned from the WLAN controller to WCS, with an average UDP packet size of 362 bytes.

Keep in mind that these traffic volumes are affected by the complexity of the controller configuration, which in turn increases the overall size of the controller configuration. The controller configuration used during this test was relatively simple; controllers with much more complex configurations generate more traffic during these operations.

WLAN Controllers and the Location Appliance

The Wireless Location Appliance learns about WLAN controllers in the management domain from WCS and when synchronized, queries these controllers using SNMP for signal strength and other information necessary to properly determine client, asset tag, and rogue location. Note that in terms of determining device location, the location appliance does not obtain device location information or signal strength data from WCS but independently polls each WLAN controller for the information it needs to perform these calculations. This occurs independently of the on-demand location capabilities present in WCS.

Wi-Fi Location Based Services: Design and Deployment Considerations discusses in detail the communication flows between the location appliance and WLAN controllers. Included in this document are traces and analysis of actual traffic flows between controllers and location appliances in both small and large footprint installations. This document is available at the following URL: http://www.cisco.com.

WCS and the Location Appliance

WCS and the location appliance exchange information regarding calibration maps and network designs during the design synchronization process. During a network design synchronization, we are generally transferring network design information from the more current to the less current partner in order to promote a common understanding of the overall design of the network and the environmental factors included in the most recent calibration maps. Lab analysis of the routine communication flows between WCS and the location appliance indicate that peak traffic flows occur during the synchronization and location server backup/restore processes.

Wi-Fi Location Based Services: Design and Deployment Considerations (see above) discusses in detail the communication flows between WCS and the location appliance as well as best practice recommendations that should be considered in deciding where the location appliance should be placed within your network.

Administering WCS

Administering Scheduled Tasks

WCS provides several pre-defined system tasks that address various areas of configuration and database backup, device status, and synchronization and statistics collection. The currently available scheduled tasks can be accessed via **Administration > Scheduled Task,** as shown in Figure 8-84. These tasks can be scheduled to run at pre-determined times of the day and with varying repetition intervals. When configured, each task can be administratively enabled or disabled. Any task can be submitted for immediate execution (including tasks that have been administratively disabled) by selecting the task check box and then selecting **Execute Now** from the command drop-down menu in the upper right-hand corner of the screen. This immediate execution capability allows the scheduled tasks feature to flexibly serve in dual roles:

- As a time-driven job scheduler, allowing you to accomplish basic system housekeeping chores
- As a method of performing on-demand functions that otherwise would not be accessible to you from within WCS (examples of this are the Database Cleanup and WCS Server Backup tasks).

Figure 8-84 Administration > Scheduled Tasks

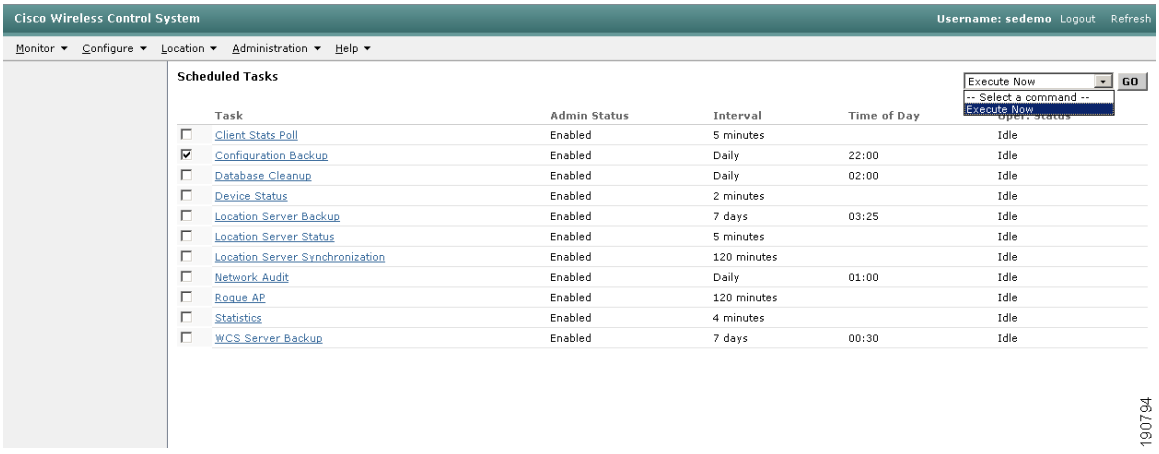

The online help system accessible via under the WCS menu bar as **Help > Online Help** contains guidance on how each of the scheduled tasks should be configured for proper operation. The following subsections provide further detail on the Configuration Backup, Network Audit, and WCS Server Backup scheduled tasks.

Configuration Backup

The Configuration Backup scheduled task archives the configurations of all controllers that have been added to WCS and are reachable at the time the task is submitted for execution.

Note For more information, see section 13.6 CSCsd54800—Config Backup Scheduled Task Fails if One Controller is Unreachable.

For each controller that is defined and currently reachable, WCS creates a configuration archive file in the default directory of the TFTP server selected with a filename that is in the format of *nnn_nnn_nnn_nnn_YYMMDD_hhmm.cfg* where *nnn* represents each octet of the IP address of the controller.

For example, a display of the tftp directory on a Linux-based WCS server shows the following files after a successful execution of this task:

```
[root@wcslinux wcs_tftp]# ls
10_1_56_10_060301_2238.cfg  10_1_56_14_060301_2237.cfg  10_1_56_18_060301_2237.cfg
10_1_56_16_060301_2238.cfg  10_1_56_12_060301_2238.cfg
[root@wcslinux wcs_tftp]#
```

To configure this task, perform the following:

Step 1 You must decide whether you desire to save and archive the current running configuration or simply archive the current saved configuration of each controller. Configuration parameters that have not been saved to the nonvolatile memory of each controller are *not* be included in the configuration archives that are produced.

a. For each controller for which you want to save and archive the current *running* configuration, go to **Configure > Controllers > Controller Properties** and ensure that the Save Before Backup check box is enabled. Click on **Save**.

   Note that this option saves the current running configuration of the controller to nonvolatile memory before archiving it. The saved configuration that was present in nonvolatile memory is overwritten. See Configuring WLAN Controllers, page 8-12, for further information on the Save Before Backup check box parameter.

b. For each controller for which you want to archive the current *saved* configuration, go to **Configure > Controllers > Controller Properties** and ensure that the Save Before Backup check box is *not* enabled. Click on **Save**.

Step 2 Go to **Administration > Scheduled Tasks** and click on the **Configuration Backup** hyperlink shown in Figure 8-84, which results in the display of the Modify Task page.

Step 3 Select the time at which you want the configuration backup task to be submitted for execution as well as the daily repetition interval.

Step 4 Choose a destination TFTP server from those provided in the drop-down list. Note that you cannot define a new TFTP server during the configuration of this task. If you want to define another server to be added to the list, use **Configure > Templates > TFTP Servers** and select Add TFTP Server from the

drop-down menu in the upper right corner of the screen. When you have added a new TFTP server in this manner, return to the Configuration Backup scheduled task panel and your newly-defined server should now be available to you.

Step 5 Enable the task by clicking on the **Admin Status Enabled** check box.

The task is now enabled for automatic submission at the time and with the daily repetition interval you have selected. If you want to schedule the task for immediate submission, select the task by enabling the check box as shown in Figure 8-84, then select **Execute Now** from the command drop-down menu in the upper right-hand screen corner and click **Go**.

Network Audit

To initiate a network audit report scheduled task, perform the following steps:

Step 1 Click **Administration > Scheduled Tasks** to display the listing of available scheduled tasks.

Step 2 Click on the hyperlink for the Network Audit scheduled task entry. This brings up the **Task > Network** Audit page. At the top of this screen the status of the last network audit is indicated. Set the time of day that you want the network audit to run and set the interval at which you want the network audit to repeat (that is, 1=daily). Finally, enable the check box to enable the network audit as shown in Figure 8-85.

Figure 8-85 Network Audit Configuration

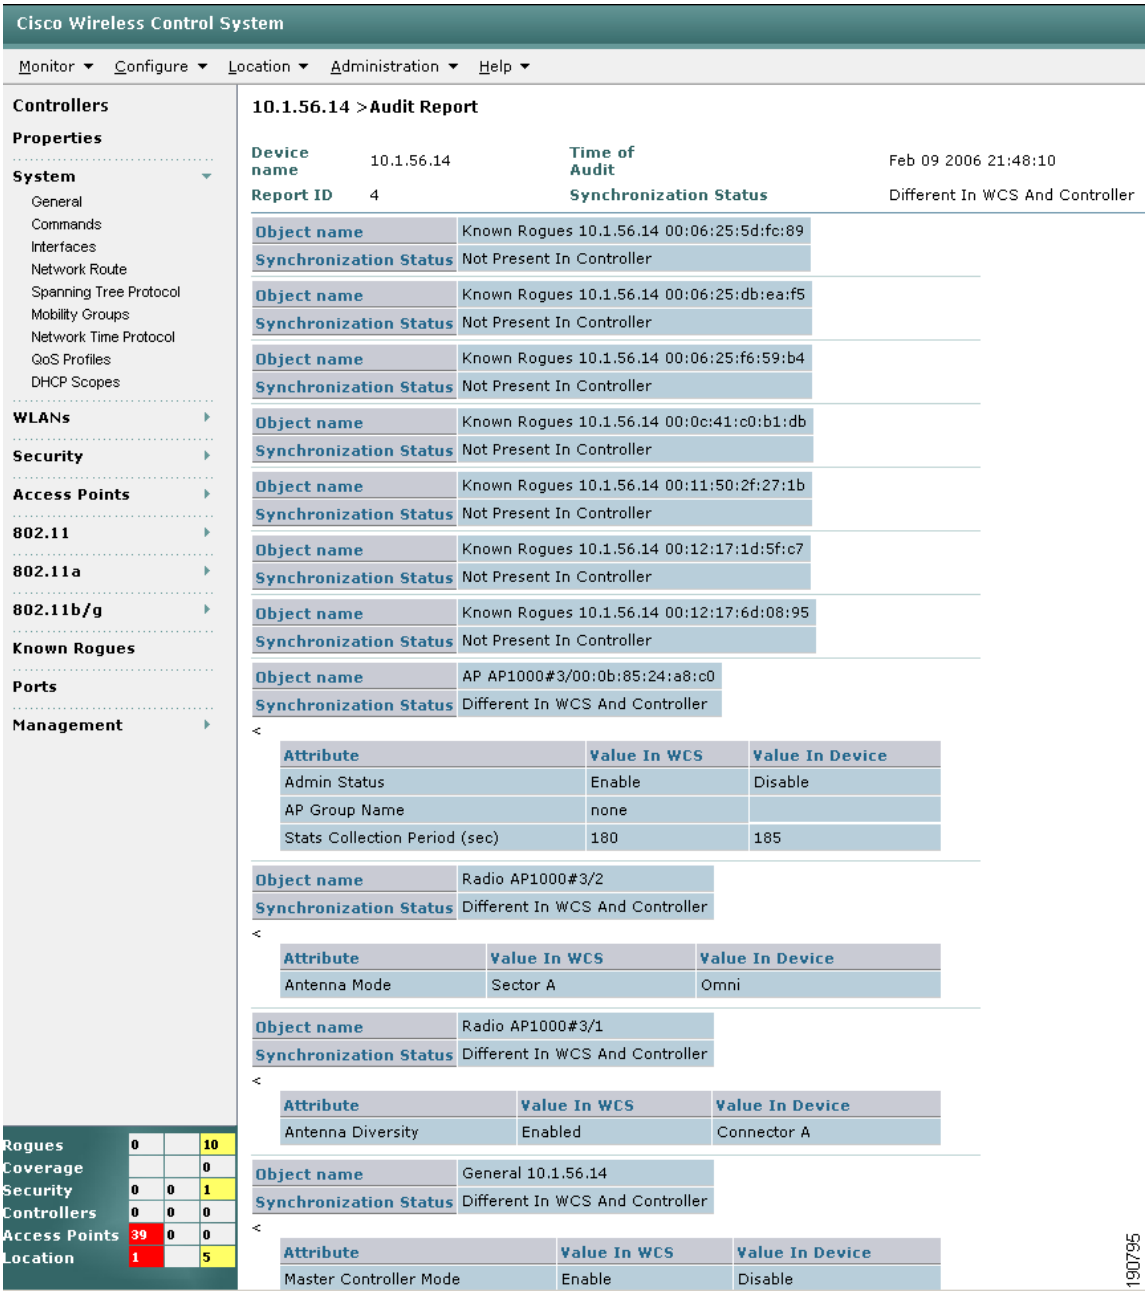

Step 3 Click **Submit**. The **Administration > Scheduled Tasks** screen should re-appear and the network audit task admin status should indicate "enabled" with the scheduled start time and repetition interval that was specified.

To view the results of a network audit that has completed running, perform the following:

Step 1 Click on **Configure > Controllers** and enable the check box of the controller for which you want to see the configuration audit report.

Step 2 Choose **View Audit Reports** from the command drop-down menu in the upper right-hand corner and click **GO**. WCS displays all available configuration audit reports for the selected WLAN controller.

Step 3 Click on the **Report ID** hyperlink of the Configuration audit report you want to view.

WCS then displays the same format configuration audit report for the WLAN controller and its registered lightweight access points as is shown in Figure 8-86.

Figure 8-86 Configuration Audit Report

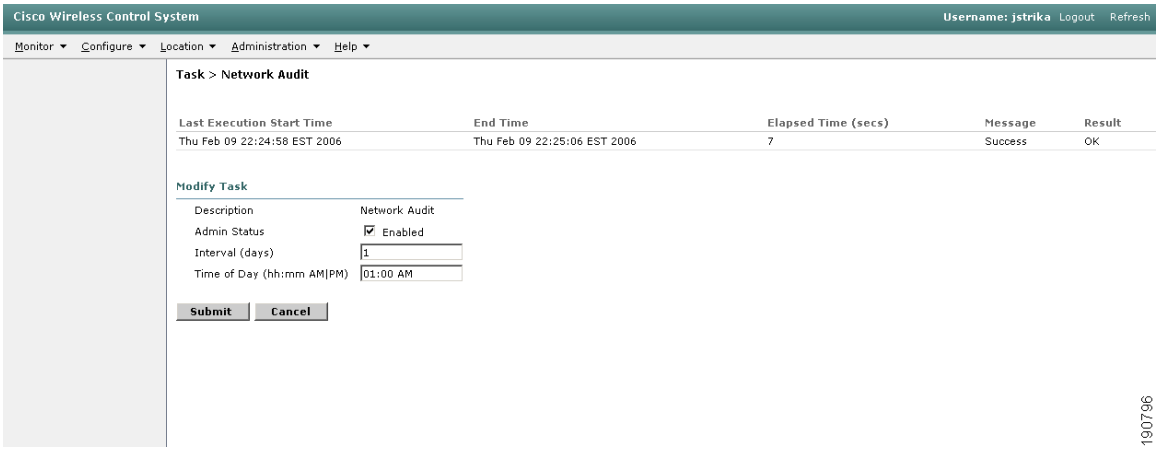

Step 4 Network Audit scheduled tasks can also submitted for immediate execution. To do this, click on **Administration > Scheduled Tasks** and then enable the check box beside the Network Audit hyperlink.

Step 5 Choose **Execute Now** from the command drop-down menu in the upper right-hand corner, and click **Go**.

WCS submits the Network Audit task for immediate execution (the status of the network audit changes to “Executing”). You may view the results of this network audit as has been previously described.

WCS Backup

The WCS Backup scheduled task provides a convenient mechanism to ensure that the WCS databases are archived on a regular basis. Unlike the “Backup” script that is executed from the operating system outside the WCS user interface, the WCS Backup scheduled task does not offer a choice of destination system or folder. Rather, the database archive is always placed on the WCS server itself in a “WCSBackup” subdirectory below the directory chosen by you at WCS installation time for the storage of FTP files. Each database archive file is named as per the following format:

DD-mon-YY_hh-mm-ss.nmsbackup.

Thus, for a WCS-Linux installation, if */opt/WCS32/wcs_ftp* was chosen as the FTP directory, the archive files created by the WCS Backup scheduled task is found in */opt/WCS32/wcs_ftp/WCSBackup*. The database archive files that can be found under that directory appears as follows:

```
[root@wcslinux WCSBackup]# ls
02-Mar-06_22-13-11.nmsbackup  02-Mar-06_22-13-47.nmsbackup  02-Mar-06_22-15-06.nmsbackup
02-Mar-06_22-13-21.nmsbackup  02-Mar-06_22-13-55.nmsbackup
02-Mar-06_22-13-30.nmsbackup  02-Mar-06_22-14-05.nmsbackup
[root@wcslinux WCSBackup]#
```

To configure the WCS Backup scheduled task, perform the following steps:

Step 1 Go to **Administration > Scheduled Tasks** and click on the **WCS Backup** hyperlink shown in the screen shown in Figure 8-84. This results in the display of the Modify Task screen.

Step 2 Specify the time at which you desire the task to be submitted for execution as well as the daily repetition interval.

Step 3 Choose the total number of database archives you want WCS to maintain on an ongoing basis (this must be at least seven and cannot exceed fifty).

Step 4 Enable the task by clicking on the **Admin Status Enabled** check box.

The task is now enabled for automatic submission at the time and with the daily repetition interval you have selected.

Step 5 If you want to schedule the task for immediate submission, select the task by enabling the check box for the task on the **Administration > Scheduled Tasks** screen, then select **Execute Now** from the command drop-down menu in the upper right-hand screen corner and click on **GO**.

Note Keep in mind that performing a WCS database backup can be relatively resource-intensive. Therefore, Cisco does not recommend that database backups be performed during peak periods of WCS usage. Schedule WCS database backups for non-busy periods when there is little use of WCS or the WCS databases.

Managing WCS Users

Adding User Accounts

WCS is installed by default with a single user *root* with a password of *public* that is a member of group *SuperUsers*. The password for *root* should be changed to a secure password as soon as possible after installation to prevent unauthorized access.

To add user accounts to WCS, use the following procedure:

Step 1 Log into WCS using the *root* account (or another account with *superuser* privileges).

Step 2 Click **Administration > Accounts** to display the All Users page.

Step 3 From the command drop-down menu in the upper right-hand corner of the page, choose **Add User** and click **GO** to display the User Administration page.

Step 4 Enter the username and password for the new WCS user account. You need to re-enter the password to confirm it.

Step 5 Under **Groups Assigned to this User**, check the appropriate box to assign the new user account to one of the six user groups supported by WCS. Keep in mind that the privileges assigned to each group can be modified further from the defaults by using the **Accounts > Groups** option described in Modifying Group Privileges, page 8-122.

a. User Assistant—Allows users only enough authority to apply an existing template to create local network user accounts for the selected controller. Local network user accounts are used to allow local controller-based authentication of clients using web authentication.

b. Lobby Ambassador—Allows an assigned user only the ability to create, apply, and delete guest user accounts that are assigned a limited lifetime of between 5 minutes and 30 days. These guest user accounts use web authentication to authenticate to the controller. For further information on the guest access capabilities of the Cisco UWN, see the "Cisco Centralized WLAN Architecture Guest Access Services" chapter in this SRND.

c. System Monitoring—Allows users to monitor WCS operations. Most general users of WCS can be assigned system monitoring capabilities. In its default configuration, it allows only general "read-only" viewing of WCS operations without the ability to affect the configuration of WCS or network components.

d. ConfigManagers—Allows users to monitor and configure WCS and network component operations. This should be assigned to users that only require the ability to change the configuration of WCS or network components managed by WCS.

e. Admin—Allows users to monitor and configure WCS operations and also perform all system administration tasks with the exception of administering WCS user accounts and passwords.

f. SuperUsers—Allows users to monitor and configure WCS operations and perform all system administration tasks including administering WCS user accounts and passwords. A SuperUser has *all* rights and privileges on the WCS system. This right should be assigned *very* judiciously.

Step 6 Click **Submit**. The name of the new user account appears on the All Users page and can be used immediately.

Modifying Group Privileges

The default privileges that are assigned to each of the groups described in the previous section are shown in Figure 8-87.

Figure 8-87 Comparison of Default WCS User Group Privileges

In some cases, it may be necessary to modify the specific privileges associated with a specific group (the procedure for accomplishing this is shown below). Keep in mind that any changes made to the group affecting the user in question *affect all users assigned to that group*.

Step 1 Click **Administration > Accounts** to display the All Users page.

Step 2 In the sidebar, click **Groups** to display the All Groups page.

Step 3 Click the name of the user group that you wish to modify. A listing of the permissions currently assigned and those available to assign for the group is displayed.

Step 4 Make any desired changes by checking or unchecking the appropriate check boxes.

Step 5 Click **Submit** to save your changes or **Cancel** to leave the settings unchanged.

Note that although users can be restricted from performing classes of activities from all lightweight access points or WLAN controllers in the WCS management domain, these privileges are not specified on a per-access point or per-controller basis.

Viewing User and Group Audit Trails

WCS allows users that are members of the *SuperUser* group to view the past WCS access audit trail of any user or group defined to WCS and also clear that audit trail if desired. User and group audit trail information is maintained indefinitely by WCS and contains the time, date, and status of authentication attempts made by each user against WCS.

To view the audit trail for a WCS user, perform the following steps:

Step 1 Log into the WCS user interface as a user assigned to the SuperUsers group.

Step 2 Click **Administration > Accounts > Users** to display the All Users page.

Step 3 Locate the user for which you want to display the audit trail information.

Step 4 Click on the  icon under the extreme right-hand column entitled “Audit Trail”. The audit trail log is displayed.

To view the audit trail for a WCS user group, use the same basic procedure but substitute **Administration > Accounts > Groups** in Step 2.

Logging Options

WCS provides extensive command logging options that are accessed and controlled via the **Administration > Logging** option from the WCS main menu bar. Logging message levels can be configured for Error, Informational, and Trace (in order of increasing detail). The default logging configuration is shown in Figure 8-88 and includes all possible logging modules enabled.

Figure 8-88 Default Logging Configuration

No.	Time	Source	Destination	SrcPort	DstPort	Protocol	Bytes	Info
173	0.011	wcswindows	AeS_4402_2	1064	snmp	SNMP	608	GET SNMPv2-SMI::ente
174	0.011	AeS_4402_2	wcswindows	snmp	1064	SNMP	640	RESPONSE SNMPv2-SMI:
175	0.022	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysU
176	0.023	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB:
177	0.024	wcswindows	AeS_4402_2	1064	snmp	SNMP	165	GET SNMPv2-SMI::ente
178	0.024	AeS_4402_2	wcswindows	snmp	1064	SNMP	178	RESPONSE SNMPv2-SMI:
179	0.031	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysU
180	0.031	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB:
181	0.032	wcswindows	AeS_4402_2	1064	snmp	SNMP	146	GET SNMPv2-SMI::ente
182	0.033	AeS_4402_2	wcswindows	snmp	1064	SNMP	152	RESPONSE SNMPv2-SMI:
183	0.040	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysU
184	0.040	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB:
185	0.042	wcswindows	AeS_4402_2	1064	snmp	SNMP	244	GET SNMPv2-SMI::ente
186	0.042	AeS_4402_2	wcswindows	snmp	1064	SNMP	255	RESPONSE SNMPv2-SMI:
187	0.050	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysU
188	0.050	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB:
189	0.053	wcswindows	AeS_4402_2	1064	snmp	SNMP	665	GET SNMPv2-SMI::ente
190	0.054	AeS_4402_2	wcswindows	snmp	1064	SNMP	700	RESPONSE SNMPv2-SMI:
191	0.063	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysU
192	0.063	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB:
193	0.064	wcswindows	AeS_4402_2	1064	snmp	SNMP	165	GET SNMPv2-SMI::ente

The standard configuration is for WCS to create up to five rotating log files with a maximum size of 2 MB each. The names are specified using the file prefixes listed (for example, wcs-0-0.log through wcs-4-0.log.) Note that a restart is required if log file size, name, or number is changed.

To view WCS log files, use the **Download** selection shown in Figure 8-88 to download the entire set of logs as a .zip compressed archive to your desktop or elsewhere for viewing with one of many popular ASCII text viewing programs, such as Microsoft Notepad, Wordpad, and so on.

There are very many log files contained in the compressed archive. However, the wcs*.log files are the first places to look when inquiring into the reason behind the abnormal termination of scheduled tasks, audit reports, configuration file archivals, and other tasks. This is also the first place to look for information about why other screen functions or commands issued from within WCS may have terminated abnormally or produced unexpected results.

When experiencing difficulties with WCS that you cannot resolve, you will likely be asked to download the compressed zip archive from the screen shown in Figure 8-88 for use by the Cisco Technical Assistance Center in resolving your problem.

Reference Publications

This chapter makes reference to the following Cisco publications:

- Cisco Wireless Control System (WCS) Installation and Upgrade Guides—
  http://www.cisco.com/en/US/products/ps6305/prod_installation_guides_list.html
- Cisco Wireless Control System Release Notes, Release 4.0—
  http://www.cisco.com/en/US/products/ps6305/prod_release_note09186a00806b0811.html
- Cisco Wireless Control System Configuration Guide, Release 4.0—
  http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_book09186a00806b57ec.html
- Deploying Cisco 44xx Series WLAN Controllers—
  http://www.cisco.com/en/US/products/ps6366/prod_technical_reference09186a00806cfa96.html

- Cisco Wireless LAN Controller Configuration Guide, Release 4.0— http://www.cisco.com/en/US/products/ps6366/products_configuration_guide_book09186a00806b0077.html
- Cisco 4400 Wireless LAN Controller—Quick Start Guide— http://www.cisco.com/en/US/products/ps6366/products_quick_start09186a00806b5e0d.html
- Quick Start Guide LWAPP-Enabled Cisco Aironet Access Points— http://cisco.com/en/US/products/hw/wireless/ps430/products_quick_start09186a00805100f5.html
- Cisco Aironet 1240AG Series Lightweight Access Point Hardware Installation Guide— http://www.cisco.com/en/US/products/ps6521/products_installation_guide_book09186a00805c90d9.html
- Release Notes for Cisco Wireless Location Appliance— http://www.cisco.com/en/US/products/ps6386/prod_release_note09186a00806b5ec7.html
- Cisco Wireless Location Appliance—Installation Guide— http://www.cisco.com/en/US/products/ps6386/products_installation_and_configuration_guide_book09186a00804fa761.html
- Cisco Wireless Location Appliance—Configuration Guide— http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_book09186a00806b5745.html
- Cisco Wireless Location Appliance—Deployment Guide— http://www.cisco.com/en/US/products/ps6386/prod_technical_reference09186a008059ce31.html
- Wi-Fi Location Based Services—Design and Deployment Considerations— http://www.cisco.com

Cisco Unified Wireless Security Integration

This chapter discusses the integration of wired network security into the Cisco Unified Wireless Solution.

Cisco provides a wide range of security features and products that are applicable to the Cisco Unified Wireless Solution. This chapter provides a collection of best practices that help integrate the most common security features and products into a wireless environment. The three areas of discussion are the following:

- Intrusion detection systems (IDS) and intrusion protection systems (IPS) integration
- Appliance and module integration
- Cisco Integrated Security Features (CISF) integration

Note A wide range of Cisco security solutions do not directly interact with the Cisco Unified Wireless Solution, but are applicable to both wired and wireless deployments. These are not discussed in detail in this design guide; for more information on security solutions, see the following URL: http://www.cisco.com/en/US/products/hw/vpndevc/index.html

IDS and IPS Integration

An IDS operates by first detecting an attack occurring at the network level, and then by either triggering a corrective action or notifying a management system so that an administrator can take action.

An IPS performs a similar analysis to that of an IDS, but is inline with the traffic flow. Rather than simply notifying network nodes about security issues, an IPS can block traffic that matches attack signatures. This gives the IPS the ability to block an attack on the fly rather than simply detecting it.

Figure 9-1 illustrates the IPS concept of the IPS being inserted into the data path, and signaling the Wireless LAN Controller (WLC) when an attack is blocked.

Figure 9-1 IPS Inline with Traffic

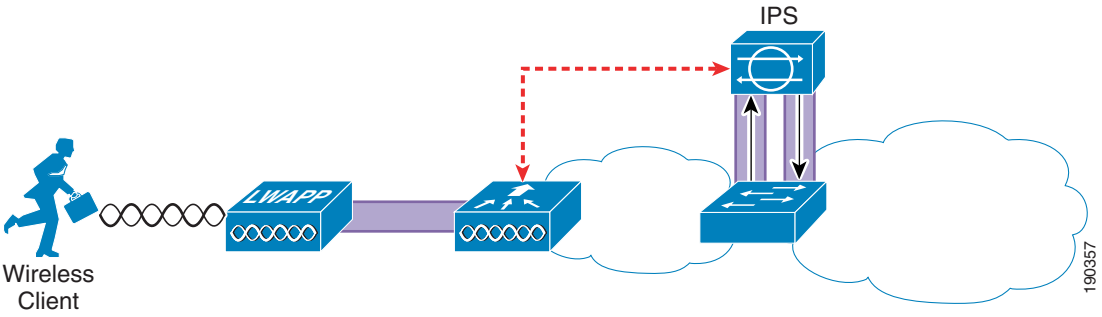

Figure 9-2 illustrates the IDS concept of the IDS analyzing; in this case, by receiving a copy of that data through a SPAN port, and signaling the WLC about an attack that should be blocked.

Figure 9-2 IDS Monitoring Traffic

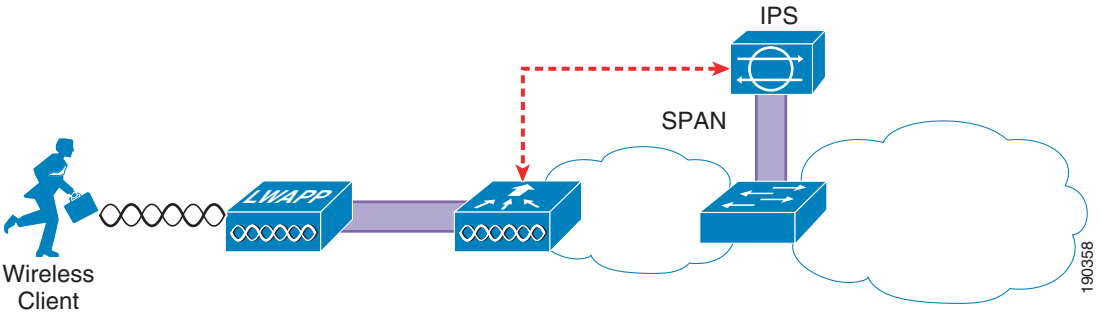

Both IDS and IPS have their place in the network. It is not the goal of this chapter to discuss the merits of either, but rather to show how a Cisco IDS or IPS can integrate with the Cisco Unified Wireless Solution to block access to the network by detected attackers.

Note For ease of reading, the term IPS is used instead of IPS/IDS for the remainder of this document.

Overview

The location of the IPS system in the network depends on the chosen security architecture; the IPS can be directly inline with the WLC, or may be in another network location, to provide protection for specific network resources. Figure 9-3 shows a schematic of the IPS deployed to protect specific network resources by analyzing all traffic to that resource. For example, the IPS may be deployed to analyze traffic to certain mission-critical servers.

Figure 9-3 Data Center IDS

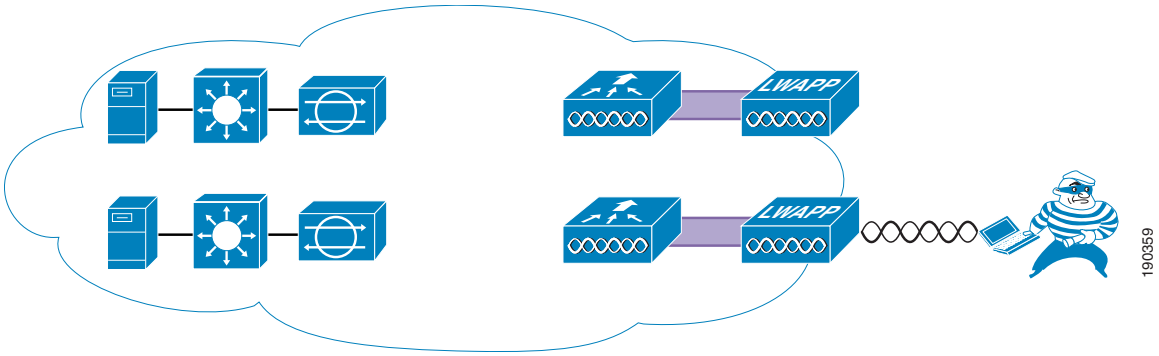

Figure 9-4 shows the IPS deployed to provide general protection of network resources by analyzing all WLAN client traffic from the WLC, ensuring that all traffic from the WLAN client is analyzed.

The chosen IDS deployment option chosen depends on the intended resource to be protected, the capacity of the IDS system, and the architecture of the WLAN system. The data center deployment has the advantage of requiring IDS hardware to be deployed *only* where the protected asset is located in the network, but does not protect against all possible WLAN exploits. The IPS-at-the-WLC deployment has the advantage of protecting against all WLAN-originated exploits, regardless of target.

Figure 9-4 IDS at the WLC

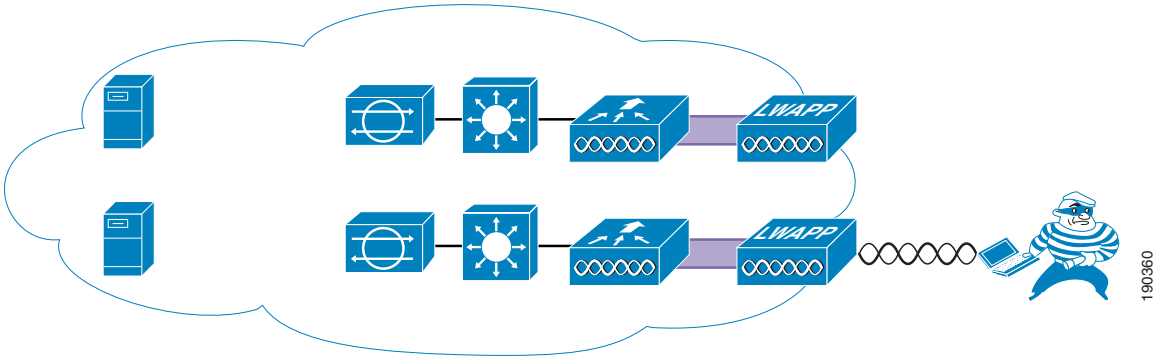

Operation

For the WLC to be able to disconnect a client attack detected by the IPS, it must learn about the client. To do this, the WLC regularly polls the IPS for information on which clients are currently targets for shunning. The IPS returns the IP addresses of currently shunned clients; the WLC uses this information to disconnect clients with those IP addresses in the WLC mobility group.

The minimum polling time is ten seconds, which means that there is a potential delay between the time that the attack is detected and the time the attacker is blocked at the WLC.

Although this means that there may be a delay in disconnecting the detected client because of the polling interval, it removes any requirement for the IPS to be aware of the network topology, and to send shun information to a specific WLC. The delay introduced by polling needs to be viewed in context of the overall IPS system, where the IPS itself has taken action to block an attack and shuns the client, and the WLC acting to augment functionality by disconnecting the offending client from the network.

Note This IPS integration features specifically uses the shunning of clients based on the information from the IPS. The decision to shun WLAN clients based on IPS information needs to be made within the context of the enterprise IPS implementation. Enterprises where the shunning of a client may cause excessive disruption to business or are vulnerable to attacks from spoofed IP address (source address checking not enforced at the access layer) may choose not to shun the client, or may only turn to shunning in special temporary circumstances. WLC features can enforce IP address spoofing protection, but the protection must extend to the remaining network to ensure that IPS shunning is not used to create a denial-of-service (DoS) attack through the spoofing of IP addresses.

WLC Configuration

Figure 9-5 shows the WLC configuration page for connection to an IPS. The WLC establishes the connection to the IPS, normally through port 443. A viewing account must exist on the IPS to which the WLC can connect through its username and password. A Transport Layer Security (TLS) certificate hash of the IPS server is also required.

The state check box controls whether the WLC attempts to connect with the IPS. If the box is checked and the configuration is correct, the connection is active; if the box is unchecked, the connection is inactive.

Figure 9-5 WLC IPS Sensor Configuration

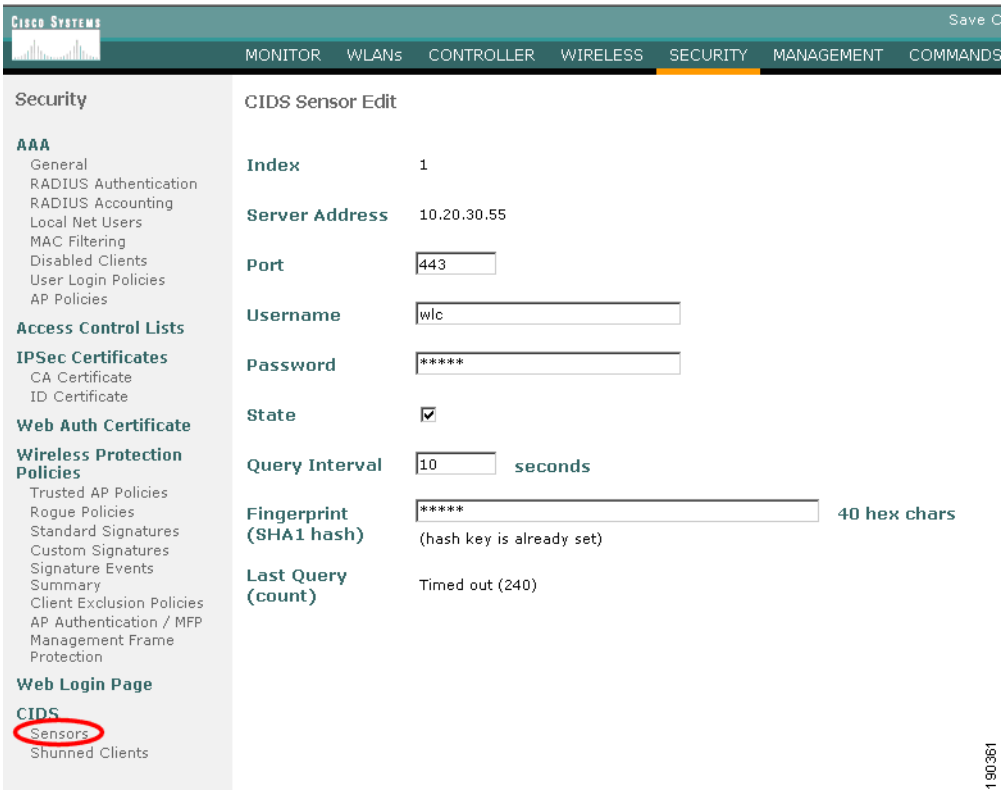

Mobility Considerations

The IPS client shun information is distributed throughout the WLC mobility group by the controller connected to the IPS. To use the information generated by an IPS, only one WLC of the mobility group needs to be a client of that IPS; all other WLCs in the mobility group receive the IPS from the connected WLC. Figure 9-6 shows a schematic of the connection between the WLC mobility group and the IPS(s). The WLC mobility group can be connected to one or more IPS(s) by any WLC member of the mobility group, and the IPS information is then distributed throughout the mobility group.

Figure 9-6 Mobility Group IDS Connection

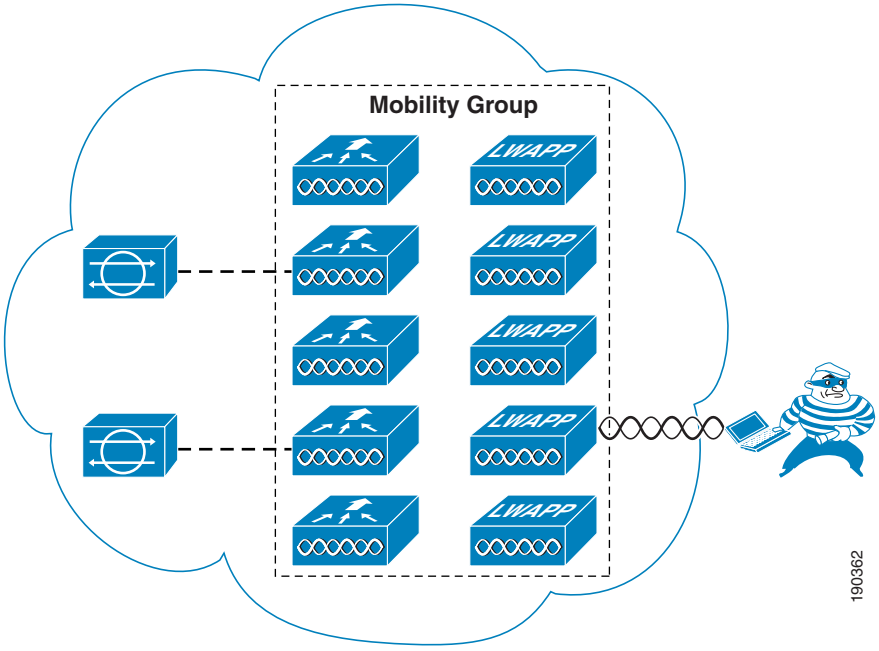

Client Shun Example

This section provides an example of a client shun on the WLC, and how it is displayed and reported.

Figure 9-7 shows the IPS report of a blocked host on the IPS.

Figure 9-7 Client Blocked at IPS

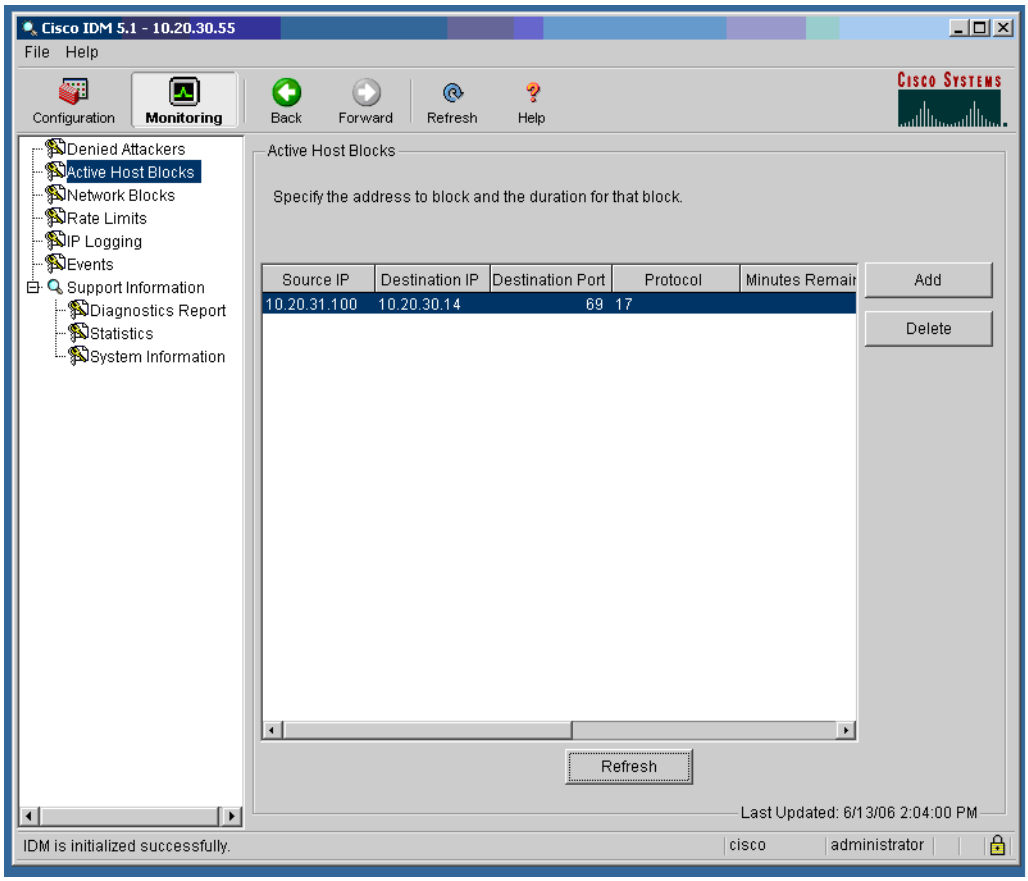

Figure 9-8 shows the shunned client report on a WLC that appears after a subsequent poll of the IPS.

Figure 9-8 Client Shun on WLC

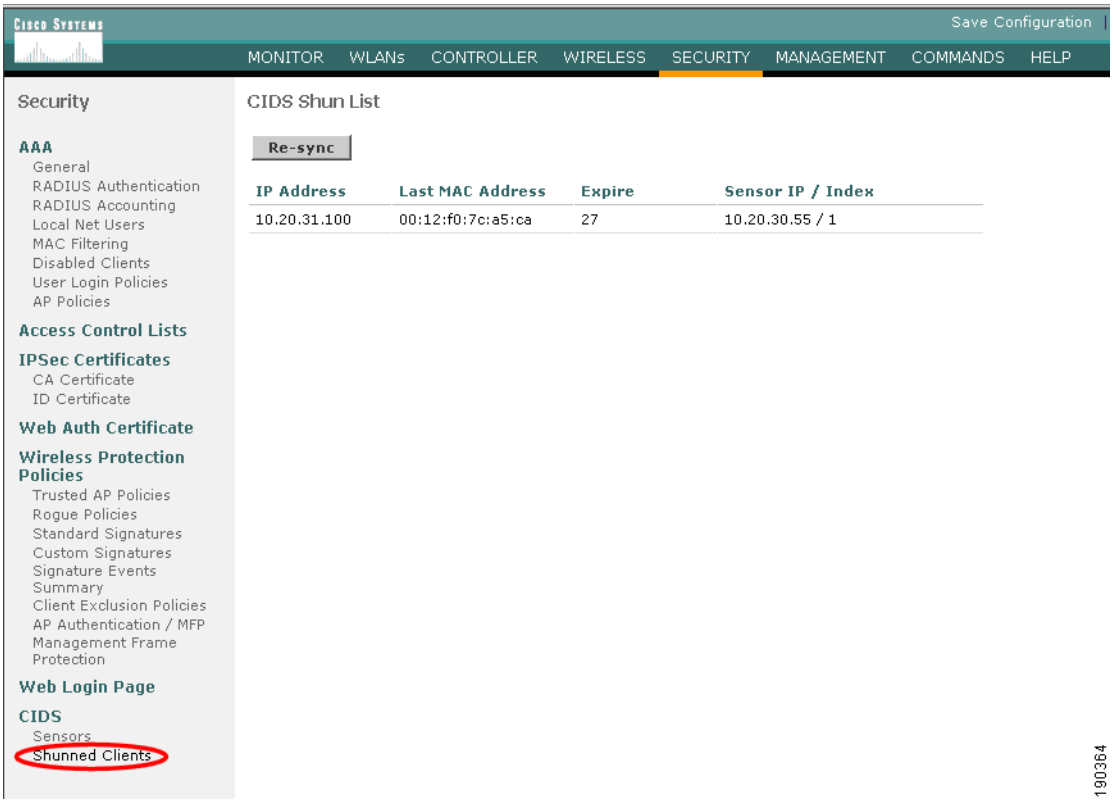

Note that the list is based on IP address. AWLC must learn the client IP address before it can learn which client MAC address to shun. Therefore, the shunning of clients based on IPS lists does not occur until the WLAN client has associated and authenticated with the WLAN. As long as the client is in the shun list on the WLCs, it is excluded; therefore, excluded clients need to be cleared from the IPS before they can cease being excluded by WLCs.

If the shunned WLAN client exists on the WLC, it is excluded, as shown in Figure 9-9.

Figure 9-9 Client Excluded

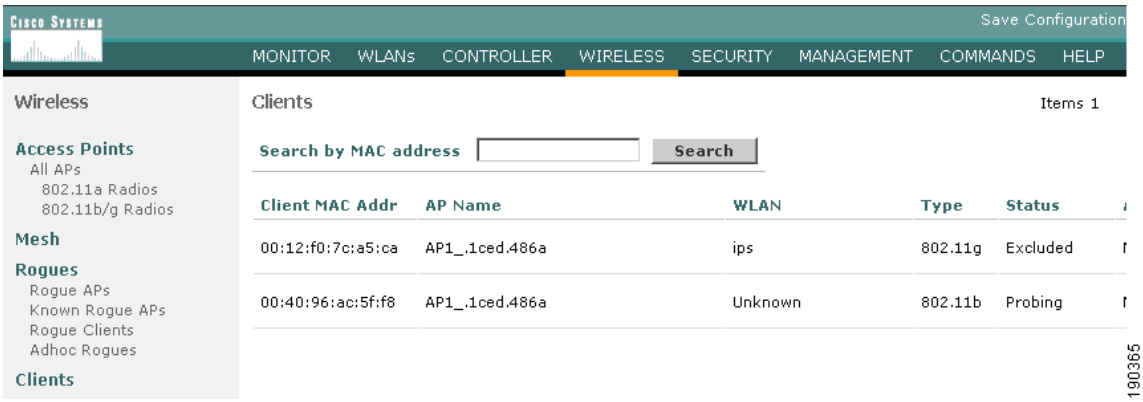

The excluding of the WLAN client is recorded on the WCS as an alarm, as shown in Figure 9-10 and Figure 9-11.

Figure 9-10 WCS Record

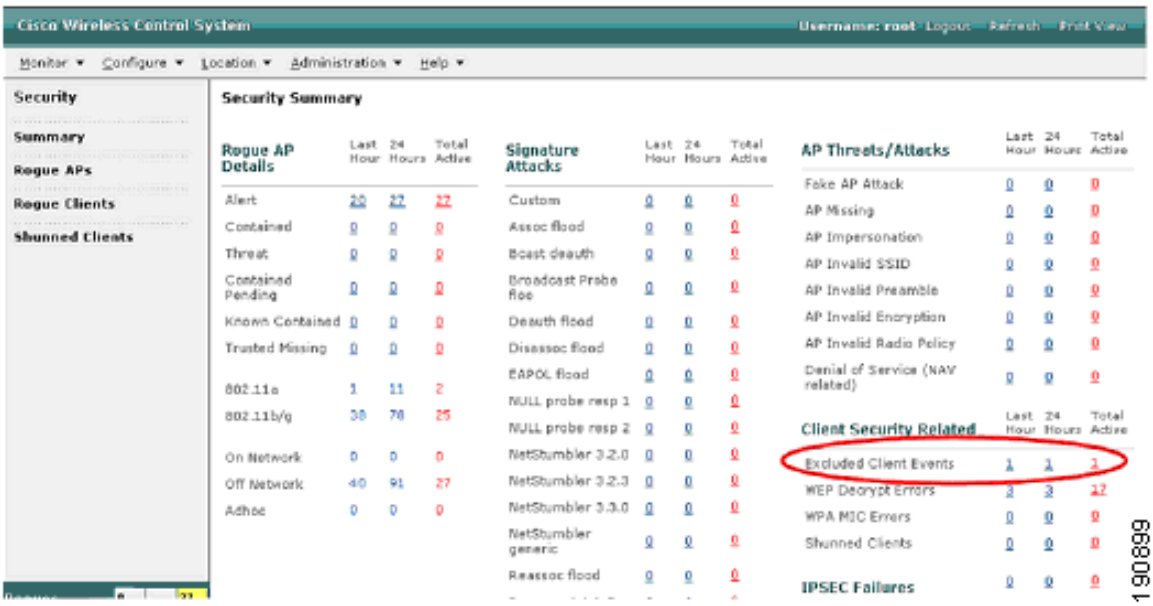

Figure 9-11 WCS Alarm Detail

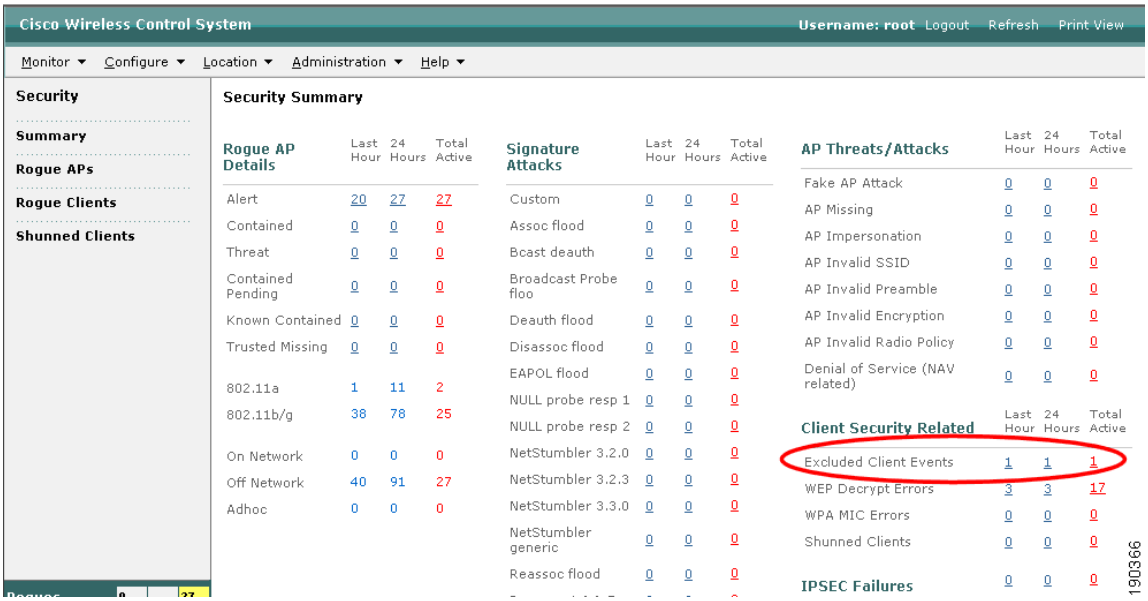

Appliance and Module Integration

Cisco provides a wide variety of security features that are either integrated into Cisco IOS, integrated into modules, or offered as appliances. The Cisco Unified Wireless architecture eases the integration of these security features into the solution because it provides a Layer 2 connection between the WLAN clients and the extended network. This means that appliances or modules that operate by being “inline”

with client traffic can be easily inserted between the WLAN clients and the core network. For example, a Cisco Wireless LAN Services Module (WLSM) implementation requires the implementation of VRF-Lite on the Cisco 6500 to ensure that the WLAN client traffic flows through a Cisco Firewall Service Module (FWSM), whereas a Cisco Wireless Services Module (WiSM) implementation can simply map the WLAN client VLAN directly to the FWSM.

The only WLAN controller not able to directly map the Layer 2 WLAN client traffic to a physical interface is the WLC ISR module (NM-AIR-WLC6). The ISR module does have access to all the IOS and IPS features available on the ISR, and the IP traffic from the WLAN clients can be forced in and out specific ISR interfaces using IOS VRF features on the router.

This section discusses the following products:

- Cisco Clean Access Server (CCAS)
- Firewall Service Module (FWSM)
- IDS Service Module (IDSM)

CCAS

The CCAS can sit between the client devices and their default gateway. This is easily achieved with the Cisco Unified Wireless Solution because the WLC provides a Layer 2 connection between the WLAN clients and the CCAS Network Admission Control (NAC) enforcement point, as shown in Figure 9-12.

Figure 9-12 Clean Access Enforcement Point

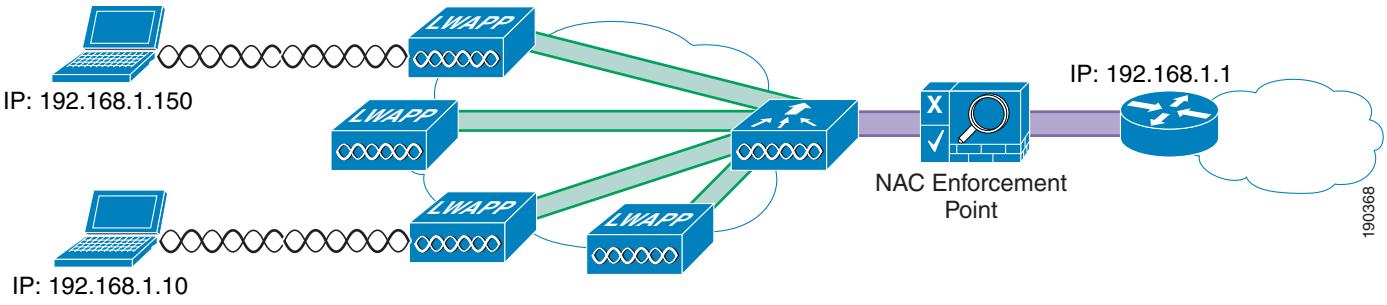

The WLAN(s) requiring NAC enforcement can be directly mapped via their VLAN to the CCAS, or these WLAN can be mapped via mobility anchors.

The schematic of Figure 9-12 is equally applicable for the integration of other security appliances into the WLAN solution where the selected WLANs are terminated directly to the untrusted interface of a security appliance.

Firewall and VPN Modules

Figure 9-13 shows the WiSM-based solution integrated with the FWSM. This allows the direct connection of selected WLANs to the FWSM.

Figure 9-13 FWSM Integration with WiSM

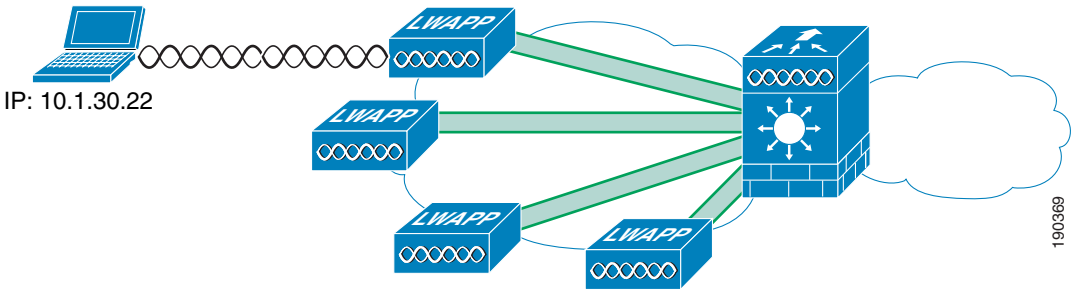

Figure 9-14 shows the logical view of the WiSM FWSM integration, where the WLAN(s) interfaces are connected to the Cisco 6500 forwarding engine through the dot1q trunk configured on the WiSM port channel interface. These WLAN interface VLANs may terminate on the 6500 routing engine or the FWSM module, or be made available outside the 6500. The FWSM interface in this case acts as the default gateway for all traffic from the selected WLAN VLANs.

Figure 9-14 FWSM Logical View

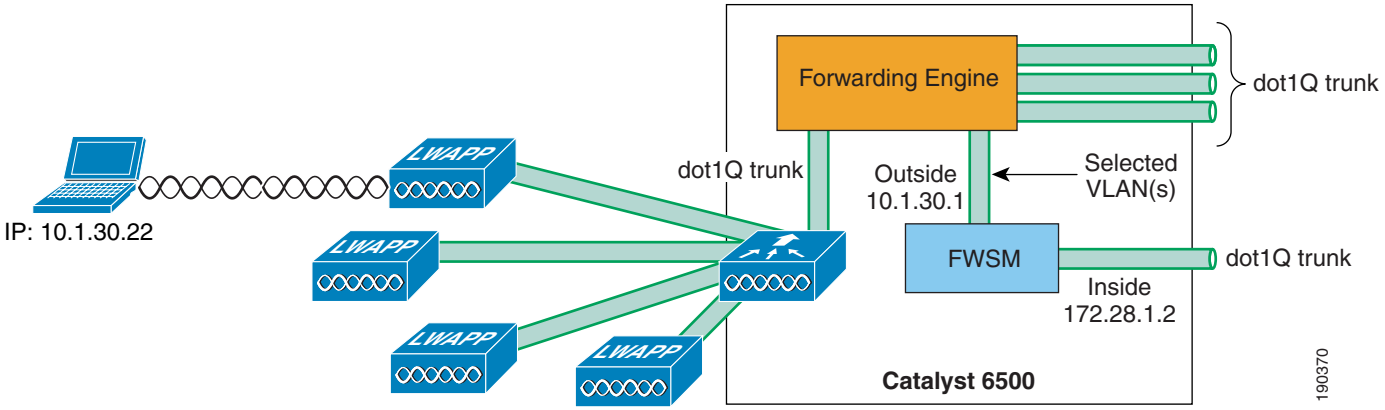

For more information on the FWSM, see the following URL:
http://www.cisco.com/en/US/products/ps6305/prod_technical_reference09186a00806053bf.html

IDSM

Figure 9-15 shows the WiSM-based solution integrated with the IDSM-2. This allows the direct connection of selected WLANs to the IDSM module through the 6500 backplane.

Figure 9-15 IDSM-2 Module

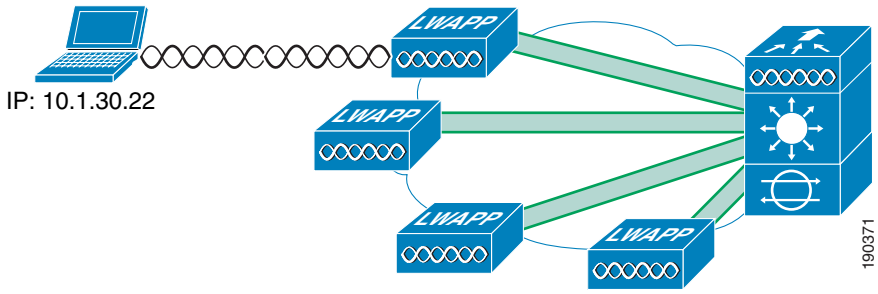

Figure 9-16 shows the logical view of the WiSM IDSM-2 integration, where the WLAN(s) interfaces are connected to the 6500 forwarding engine through the dot1q trunk configured on the WiSM port channel interface.

Figure 9-16 IDSM-2 Logical View

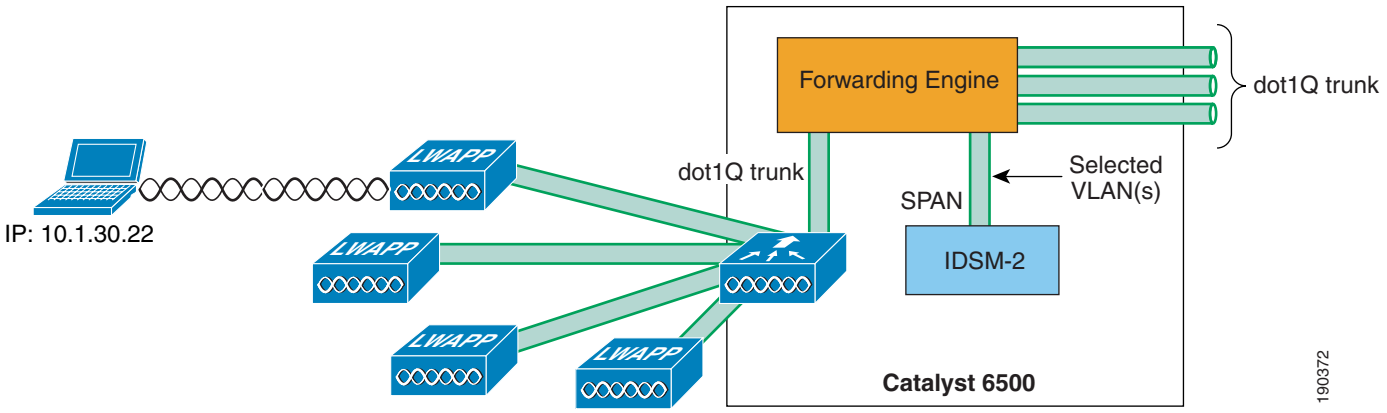

These WLAN interfaces VLANs may terminate on the 6500 routing engine or be made available outside the 6500, and can also be "spanned" to the IDSM-2. The IDSM-2 is in a passive mode, monitoring traffic to and from the selected WLAN interfaces.

The IDSM-2 may also be connected inline with the WLAN traffic in a similar manner to the FWSM if the module is being implemented for IPS purposes.

For more information on the IDSM-2, see the following URL:
http://www.cisco.com/en/US/partner/products/hw/modules/ps2706/products_data_sheet09186a00801e55dd.html

Cisco Integrated Security Features Integration

Cisco Integrated Security Features (CISF), are available on Cisco Catalyst switches, and help mitigate against a variety of attacks that a malicious user might launch after gaining wireless access to the network. This section describes these attacks, how a WLC protects against these attacks, and how CISF, when enabled on the access switch, can help protect the network. It includes the following topics:

- Overview
- Attacks the CISF can help prevent

- CISF for wireless test scenarios
- CISF for wireless test results
- Summary of test results
- Conclusion

Note This document describes only the attacks that CISF can help prevent when enabled on access switches, and is not meant to be a comprehensive analysis of all the possible attacks that are possible on wireless networks.

Overview

Attacks can occur against either wired or wireless networks. However, a wireless network connection allows an attacker to craft an attack without physical access to the network. The WLC and CISF include features that are specifically designed to prevent attacks, including the following:

- MAC flooding attacks
- DHCP rogue server attacks
- DHCP exhaustion attacks
  - ARP spoofing attacks
  - IP spoofing attacks

MAC Flooding Attack

MAC flooding attacks are attempts to fill a switch Content-Addressable Memory (CAM) table and force the switch to start flooding LAN traffic. These attacks are performed with tools such as *macof* (part of the dsniff package), which generates a flood of frames with random MAC and IP source and destination addresses.

The Layer 2 learning mechanism of an Ethernet switch is based on the source MAC addresses of packets. For each new source MAC address received on a port, the switch creates a CAM table entry for that port for the VLAN to which the port belongs. The macof utility typically fills the CAM table in less than ten seconds, given the finite memory available to store these entries on the switch. CAM tables are limited in size. If enough entries are entered into the CAM table before other entries expire, the CAM table fills up to the point that no new entries can be accepted.

A network intruder can flood the switch with a large number of bogus-source MAC addresses until the CAM table fills up. The switch then floods all its ports with incoming traffic because it cannot find the port number for a particular MAC address in the CAM table. The switch, in essence, acts like a hub to the detriment of performance and security. The overflow floods traffic within the local VLAN, so the intruder sees traffic within the VLAN to which he or she is connected.

At Layer 3, the random IP destinations targeted by macof also use the multicast address space. Thus, the distribution layer switches that have multicast turned on experience high CPU usage levels as the protocol independent multicast (PIM) process attempts to handle the false routes.

DHCP Rogue Server Attack

The DHCP rogue server event may be the result of an attack, or a user may accidentally bring up a DHCP server on a network segment and begin inadvertently issuing IP addresses. An intruder may bring up a DHCP server to issue an address with DNS server or default gateway information that redirects traffic to a computer under the control of the intruder.

DHCP Starvation Attack

DHCP starvation attacks are designed to deplete all of the addresses within the DHCP scope on a particular segment. Subsequently, a legitimate user is denied when requesting a DHCP IP address and thus is not able to access the network. *Gobbler* is a public domain hacking tool that performs automated DHCP starvation attacks. DHCP starvation may be purely a DoS mechanism or may be used in conjunction with a malicious rogue server attack to redirect traffic to a malicious computer ready to intercept traffic.

ARP Spoofing-based Man-In-the-Middle Attack

A man-in-the-middle (MIM) attack is a network security breach in which a malicious user intercepts (and possibly alters) data traveling along a network. One MIM attack uses ARP spoofing. ARP spoofing is a technique in which a gratuitous Address Resolution Protocol (ARP) request is used to misdirect traffic to a malicious computer such that the computer becomes the “man in the middle” of IP sessions on a particular LAN segment. The hacking tools *ettercap*, *dsniff*, and *arpspoof* may be used to perform ARP spoofing. Ettercap in particular provides a sophisticated user interface that displays all the stations on a particular LAN segment and includes built-in intelligent packet capturing to capture passwords on a variety of IP session types.

IP Spoofing Attack

IP spoofing attacks spoof the IP address of another user to perform DoS attacks. For example, an attacker can ping a third-party system while sourcing the IP address of the second party under attack. The ping response is directed to the second party from the third-party system.

CISF for Wireless

This section describes the various unified wireless deployment scenarios used and how the WLC or the CISF features defend against wireless attacks.

CISF is currently available only on the access switch, not directly on the access point (AP); thus, the benefits of these features are available only if the traffic from the wireless attacker goes through the switch.

The definition of an access switch is slightly different in the Unified Wireless solution, because three locations can be considered an access switch:

- The point that a controller interface terminates on the network
- The point that a standard Lightweight Access Point Protocol (LWAPP) AP terminates on the network
- The point that a hybrid remote edge access point (H-REAP) terminates on the network

These locations are illustrated in Figure 9-17.

Figure 9-17 Access Switches

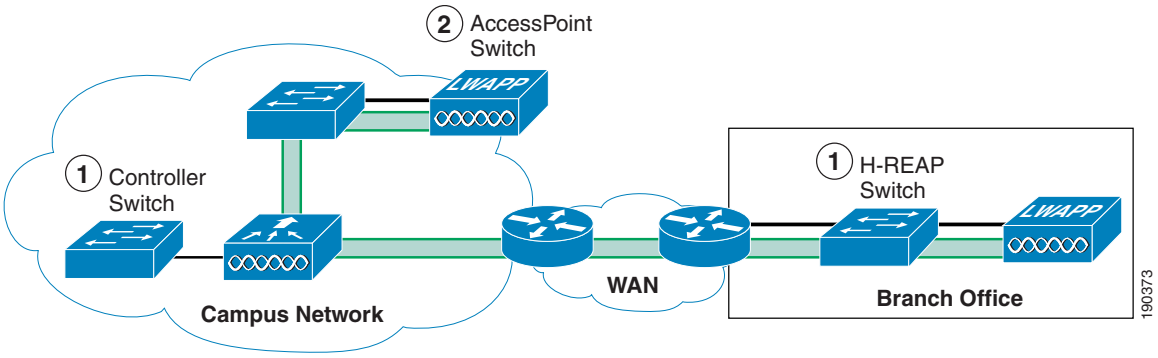

The connections of interest in the CISF discussions are the controller switch and the H-REAP switch. The AP switch is not discussed because WLAN traffic does not terminate on this switch, and the AP simply appears as a single device connected to that switch port, so from a security point of view it can be considered an access client.

Note The primary difference between the LWAPP AP and a standard client is that the differentiated services code point (DSCP) of the LWAPP AP should be trusted.

The scope of these investigations is limited to attacks between wireless users, because of the recommended design guidance that wireless and wired users should be kept on separate subnets, and discussing attacks across subnet boundaries is beyond the scope of this discussion.

The three following scenarios are considered:

- Scenario 1—Target is associated to the same AP to which the attacker is connected
- Scenario 2—Target is associated to a different AP than the attacker
- Scenario 3—Target is associated to a different AP than the attacker, and this AP is connected to a different controller

For Scenario 1, in which both attacker and target are associated to the same AP, the traffic remains local to the H-REAP or WLC, and CISF is not useful. In this case, explore other alternatives are explored to mitigate the effects of the attacks. The second and third scenarios are the ones in which CISF can be effective.

For enterprise WLAN deployment, Cisco recommends the use of multiple VLANs per SSID. This requires configuring an 802.1q trunk between the Fast Ethernet port on the AP and the corresponding port on the access switch. With multiple VLANs defined, the administrator can keep the data traffic separated from the AP management traffic. The company security policy is also likely to require having different types of authentication and encryptions for different type of users (open authentication and no encryption for guest access, dot1x authentication and strong encryption for employees, and so on). This is achieved by defining multiple SSIDs and VLANs on the AP.

Given the above, the example configurations use a trunk connection between the WLC or H-REAP AP and the access switch.

CISF for Wireless Application

This section describes each of the features provided within CISF that were tested for protection against wireless attacks, and includes the following topics:

- Using port security to mitigate a MAC flooding attack
- Using port security to mitigate a DHCP starvation attack
- Using DHCP snooping to mitigate a rogue DHCP attack
- Using Dynamic ARP Inspection to mitigate a man-in-the-middle attack
- Using IP Source Guard to mitigate IP and MAC spoofing

Using Port Security to Mitigate a MAC Flooding Attack

This section describes how to use CISF port security to mitigate a MAC flooding attack. It includes the following topics:

- Port security overview
- Port security in a wireless network
- Effectiveness of port security

Port Security Overview

Port security sets a maximum number of MAC addresses allowed on a port. You can add addresses to the address table manually, dynamically, or by a combination of the two. Packets are dropped in hardware when the maximum number of MAC addresses in the address table is reached, and a station that does not have a MAC address in the address table attempts to send traffic.

Enabling port security on the access port of the switch stops a MAC flooding attack from occurring because it limits the MAC addresses allowed through that port. If the response to a violation is set to **shutdown**, the port goes to error-disable state. If the response is set to **restrict**, traffic with unknown source MAC addresses are dropped.

Port Security in a Wireless Network

It is not generally recommended to enable port security on a switch port connected to an H-REAP AP or WLC. The use of port security implies knowing the exact number of MAC addresses that the switch learns and allows from that port; in the case of an H-REAP AP or WLC, the various source MAC addresses that the switch learns usually correspond to wireless users. Setting port security on the switch port allows only a certain number of users on the wired network.

For example, a company might have a security policy that allows only certain MACs, and a certain number of them, to send traffic through the access point. In this case, a combination of MAC filtering on the H-REAP AP or WLC and port security on the switch ensures that only the selected users access the wired network. Most of the time, however, a company implements a WLAN to facilitate the mobility of the employees, which implies that an H-REAP AP or WLC, at any given time, does not have a predetermined number of users associated with it. In all cases in which it is impossible to determine the number of users connected to the AP, enabling port security on the switch port brings no advantages. At worst, it can create an involuntary DoS attack; if the policy for port security is set to shut down the port in the case of a violation, when this happens, all the users connected to that AP lose network connectivity. Figure 9-18 shows an example of using port security to limit a wireless MAC flooding attack by locking down the port and sending an SNMP trap.

Figure 9-18 Using Port Security

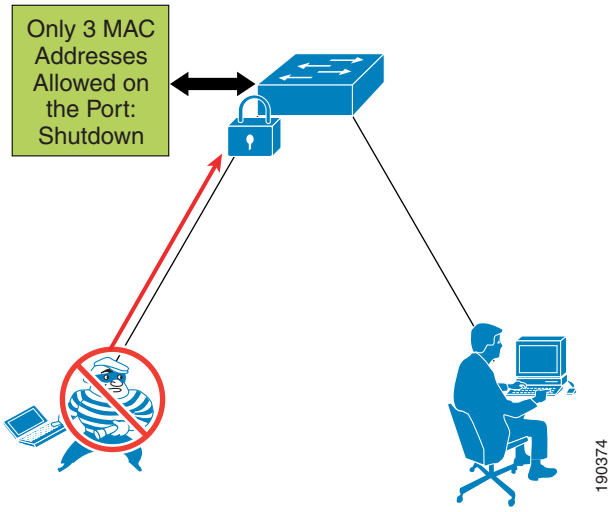

Effectiveness of Port Security

Even if port security is not an option to stop this attack (as explained), the MAC flooding attack is unsuccessful when launched by a wireless user. The reason for this is the 802.11 protocol itself. The association to an AP is MAC-based; this means that the AP bridges (translational bridge) traffic coming only from or going to known users or known MACs. If a MAC flooding attack is launched from a wireless user, all the 802.11 frames with random source MAC addresses that are not associated to the AP are dropped. The only frame allowed is the one with the MAC of the malicious user, which the switch has probably already learned. Thus, the operation of the access point prevents the switch from being susceptible to MAC flooding attacks.

Using Port Security to Mitigate a DHCP Starvation Attack

This section describes how to use CISF port security to help prevent a DHCP starvation attack. It includes the following topics:

- Overview
- Wireless DHCP starvation attack
- Effectiveness of port security

Overview

For wired access, port security can currently prevent a DHCP starvation attack launched from a PC connected to a switch and using a tool such as Gobbler. The failure of the attack is due more to a limitation of the tool than an actual fix provided by port security. The only reason such an attack fails is that

Gobbler uses a different source MAC address to generate a different DHCP request; if the attacker used his or her MAC address in the Ethernet packet and simply changed the MAC address in the DHCP payload (the field is called chaddr), port security would not stop the attack.

All that can currently be done is to try to slow down the attack using a DHCP rate limiter on the switch port. The next software release for the Catalyst switches will provide the fix for such attack: the switch will need to compare the source MAC address of the DHCP request with the MAC address in the DHCP payload and drop the request if the two are different. This fix assumes that the DHCP server is connected to the wired infrastructure, so it will be useful if the DHCP server feature available on the AP is used.

Wireless DHCP Starvation Attack

In a Unified Wireless deployment, the vulnerability to a DHCP starvation attack depends on whether the WLC terminates the user traffic or an H-REAP terminates the user traffic.

The WLC protects the network from DHCP starvation attacks because it examines DHCP requests to ensure that the client MAC address matches the chaddr. If the addresses do not match, the DHCP request is dropped.

In case the H-REAP VLAN is terminated locally, the DHCP request does not go through the controller and an analysis of the chaddr cannot performed. In this case, the same considerations related to this attack for wired access also apply when the attacker launches the attack via wireless. A smart attacker uses the MAC address with which he or she is associated to the AP to generate the random DHCP requests, simply changing the requesting MAC address in the DHCP packet payload. To the AP, the packet looks like a valid packet coming in from one associated client and it does not drop it.

Using DHCP Snooping to Mitigate a Rogue DHCP Server Attack

This section describes how to use DHCP snooping to mitigate a DHCP server attack. It includes the following topics:

- Overview
- DHCP snooping for wireless access
- Effectiveness of DHCP snooping

Overview

DHCP snooping is a DHCP security feature that provides security by building and maintaining a DHCP snooping binding table and filtering untrusted DHCP messages. It does this by differentiating between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCP server or another switch. End-user ports can be restricted to sending only DHCP requests and no other type of DHCP traffic. Trusted ports allow any DHCP message to be forwarded. The DHCP snooping table is built per VLAN and ties the IP address/MAC address of the client to the untrusted port. Enabling DHCP snooping prevents users from connecting a non-authorized DHCP server to an untrusted (user-facing) port and start replying to DHCP requests.

DHCP Snooping for Wireless Access

The WLC manages all DHCP requests from clients and acts as a DHCP relay agent. DHCP requests from WLAN clients are not broadcast back onto the WLAN, and they unicast from the WLC to the configured DHCP server. This protects the WLAN clients connected to the WLC from rogue DHCP server attacks.

Clients connected to an H-REAP native interfaces are not protected against rogue DHCP server attacks by the WLC.

Keep in mind that the CISF features (in this case DHCP snooping) are available on the switch and not on the AP, so they intercept the malicious messages only if the traffic from the rogue server goes through the switch.

Figure 9-19 shows an example of using DHCP snooping to mitigate against a rogue DHCP server attack, and how the attack can happen before the switch provides DHCP protection.

Figure 9-19 Security Used Against Rogue DHCP Server Attack

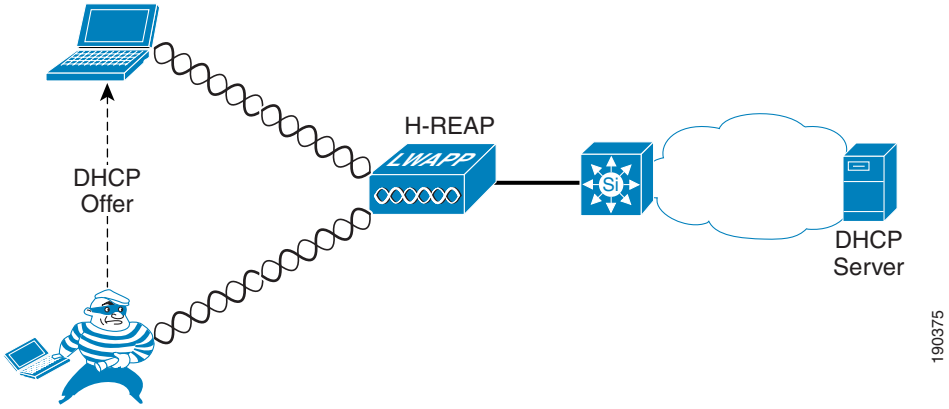

Effectiveness of DHCP Snooping

DHCP snooping is enabled on a per-VLAN base, so it works on a trunk port. A separate DHCP snooping entry is inserted for each DHCP request received on the same trunk port for clients in different VLANs. The fact that DHCP snooping works for a trunk port is very important because it makes this CISF feature applicable to a WLAN deployment in which multiple SSIDs/VLANs are requested on the H-REAP native interfaces. If the DHCP snooping attacker is associated to a different H-REAP, the switch is able to protect against the attack. However, if the attacker and the target are associated to the same H-REAP, the attack does not go through the switch and the attack is not detected.

DHCP snooping also provides some protection against DHCP server attacks by rate limiting the DHCP requests to the DHCP server.

Using Dynamic ARP Inspection to Mitigate a Man-in-the-Middle Attack

This section describes how to use Dynamic ARP Inspection (DAI) to mitigate ARP spoofing MIM attacks. It includes the following topics:

- Overview
- DAI for wireless access
- Effectiveness of DAI

Overview

DAI is enabled on the access switch on a per-VLAN basis and compares ARP requests and responses, including gratuitous ARPs (GARPs), to the MAC/IP entries populated by DHCP snooping in the DHCP binding table. If the switch receives an ARP message with no matching entry in the DHCP binding table, the packet is discarded and a log message is sent to the console. DAI prevents ARP poisoning attacks that may lead to MIM attacks such as those launched using ettercap (an example of a tool that has a very intuitive user interface) by stopping the GARP messages that the malicious user sends to the target to alter their ARP table and receive their traffic. The ARP messages are filtered directly at the port to which the attacker is connected.

DAI for Wireless Access

The WLC protects against MIM attacks by performing a similar function to DAI on the WLC itself. This allows it to block the GARPs necessary for this attack to proceed. DIA should not be enabled on the WLANs behind the WLCs, because the WLC uses the GARP in Layer 3 roaming for clients.

For the H-REAP, two different scenarios can impact the effectiveness of the DAI feature on the switch. It is possible to have DAI enabled on all the VLANs carried on the trunk from the AP to the switch. This makes the CISF feature applicable in a wireless environment in which multiple SSIDs/VLANs are deployed on the AP. The following scenarios assume the attacker to be associated to an AP and Layer 2 adjacent to the targets:

- Scenario 1—One of the targets is wireless and associated to the same AP as the attacker (the other target is the default gateway). This is considered the most typical attack.
- Scenario 2—The targets are both wireless.

These different scenarios determine in which cases the traffic goes through the switch and thus can be stopped.

In Scenario 1, the MIM attack attempts to use a GARP to change the ARP table entries on the default gateway and the wireless victims, redirecting traffic to go through the attacker. DAI can block the GARP at the default gateway, but has no impact on the GARP to the wireless client. This limits the effectiveness of the MIM attack.

In Scenario 2, the MIM attack sends GARPs to wireless clients, and the switch implementing DAI does not see these GARPs and cannot block the attack.

Figure 9-20 shows an example of the attack mechanism where GARPs are sent to the two IP connection nodes on the subnet to divert the traffic between them.

Figure 9-20 Dynamic ARP Inspection

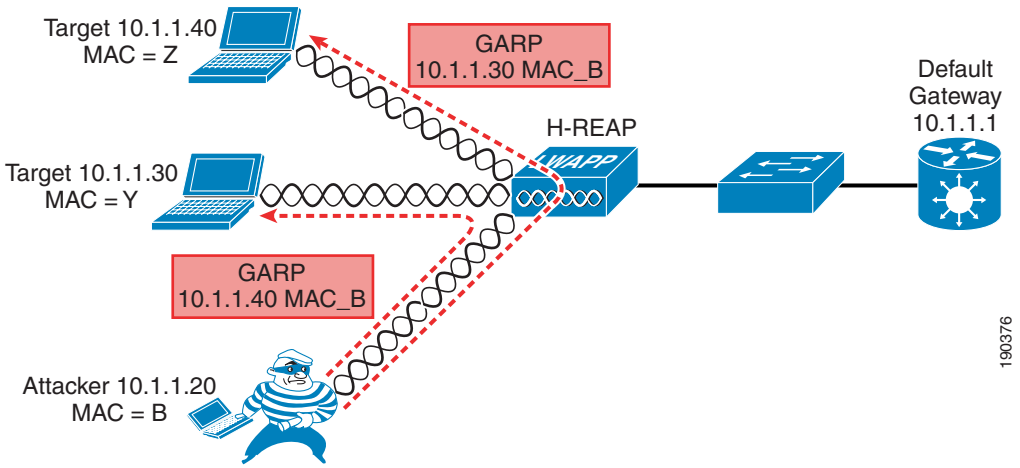

Effectiveness of DAI

The attack is completely successful only in the example of Figure 9-20 in which the traffic remains local to the H-REAP and never goes through the switch. Usually the interesting traffic for an attacker, such as passwords and account information, travels from the wireless client to the wired network (server or Internet), so this is not too harmful.

The scenario where the default gateway and a wireless client are the attack targets can be called a half duplex MIM attack. Ettercap is able to modify the ARP table of the wireless user that is now sending all the traffic to the intruder, but the GARP to the default gateway is intercepted by the switch and a message is logged, as shown in the following example:

```
4507-ESE#sh ip arp inspection log
Total Log Buffer Size : 32
Syslog rate : 5 entries per 1 seconds.
Interface Vlan Sender MAC Sender IP Num of Pkts Reason
---------- ---- -------------- --------------- ----------- ------
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Wed Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:48 PDT Tue Feb 3 2003) DHCP
Deny
Interface Vlan Sender MAC Sender IP Num of Pkts Reason
---------- ---- -------------- --------------- ----------- ------
Fa3/26 20 00d0.5937.7acc 10.20.1.100 1(11:07:49 PDT Tue Feb 3 2003) DHCP Deny
```

Because the MAC address is provided in the log, the administrator can take further action to block the attack by disassociating the attacker.

When DAI is configured on the VLAN, an ARP rate limiter is configured globally to prevent flooding of ARP requests coming from a certain port. The default value of the rate limiter is 15 packets per second (pps). If this limit is reached, the switch disables the port to prevent the attack. In this case, to launch a MIM attack, an attacker must first discover who else is Layer 2 adjacent. To do this, ettercap generates a series of GARPs, claiming to be each one of the IP address on the subnet. In this way, the real owner of that address replies and ettercap can build its table.

In lab tests, this limit has been reached immediately by ettercap and the port shut down. This is acceptable in a wired scenario, but in a wireless scenario, by shutting down the port connected to the AP, all the wireless users lose their connection to the outside world and a possible MIM attack turns into a DoS attack.

To avoid this involuntary attack created by enabling DAI, Cisco recommends turning off the ARP rate limiter on the port of the switch connected to the AP. You can do this with the following interface level command:

```
ip arp inspection limit none
```

An alternative is to change the threshold value to a value larger than 15 pps. However, this is not a general remedy because it depends on the implementation of the specific tool being used to launch the attack.

Using IP Source Guard to Mitigate IP and MAC Spoofing

This section describes how to use IP Source Guard to mitigate IP and MAC spoofing. It includes the following topics:

- Overview
- IP Source Guard for wireless access
- Effectiveness of IP Source Guard

Overview

When enabled on an interface of the access switch, IP Source Guard dynamically creates a per-port access control list (PACL) based on the contents of the DHCP snooping binding table. This PACL enforces traffic to be sourced from the IP address issued at DHCP binding time and prevents any traffic from being forwarded by other spoofed addresses. This also prevents an attacker from impersonating a valid address by either manually changing the address or running a program designed to do address spoofing, such as hping2. This feature has an option (port security) to filter the incoming address, also using the MAC address in the DHCP snooping binding table.

The attacker typically uses the spoofed address to hide his or her real identity and launch an attack, such as a DoS attack, against a target.

IP Source Guard for Wireless Access

In the case of wireless access, IP Source Guard can be enabled on the trunk port connecting the access switch to the IP. This allows the switch to filter any traffic coming from wireless users that does not match an entry in the DHCP binding table.

IP Source Guard does not need to be enabled on the VLANs behind a WLC, because the WLC performs a similar function to ensure that the IP address used by a client is the IP address that has been assigned to that client.

IP Source Guard is applicable to the H-REAP because it does not maintain a check on WLAN client MAC address IP address bindings.

In tests, the following two scenarios were considered:

- Scenario 1—The target is represented by another wireless user associated to the same AP.
- Scenario 2—The target is another wireless user associated to a different AP.

Figure 9-21 shows an example of using IP Source Guard to mitigate IP and MAC spoofing attacks.

Figure 9-21 IP Source Guard Preventing MIM

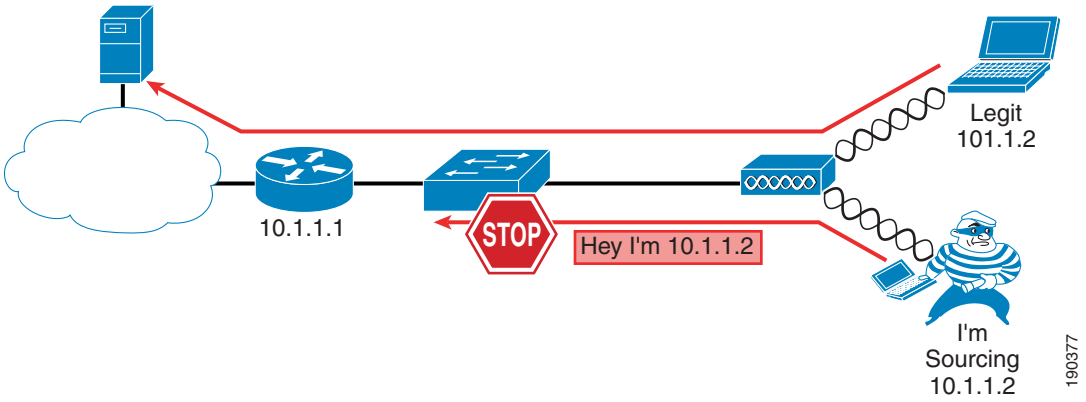

Effectiveness of IP Source Guard

The effectiveness of this feature depends on two factors: the way the attacker is able to spoof the address, and which scenario is being tested.

The association to the AP is based on the MAC address, so if the AP receives a frame with an unknown source address, it drops the frame. When launching an IP spoofing attack, the attacker has the option to use his or her own MAC address or to use one from another user connected to the same AP. All the other combinations, such as using a random MAC address or using the MAC address of a user connected to another AP, lead to a failed attack because the AP drops the frame.

In case the attacker uses his or her own MAC address but spoofs the IP address, IP Source Guard enabled on the switch stops the attack in all the scenarios described above except the first. In the first scenario, the traffic stays local to the AP and the CISF feature never kicks in. In the other scenarios, CISF successfully stops the attack because the IP-spoofed packet sent by the malicious user has no entry in the DHCP snooping table.

However, if the attacker is able to spoof both the MAC and the IP address of another wireless user connected to the same AP, basically assuming the identity of another user, then the attack is successful even in Scenarios 2 and 3.

Spoofing both the Mac and IP address is realistically possible in a hot spot environment where no encryption is used, or when the weaknesses of Wired Equivalent Privacy (WEP) are exploited. It has been shown in tests that you need to passively listen to only six million packets to break the encryption mechanism used by WEP.

This is one of the reasons why Cisco highly recommends the use of strong encryption whenever possible, preferably dynamic keys with TKIP and MIC, to make it harder for the attacker to break the encryption and use the key, IP address, and MAC address of another user to launch an attack.

Summary of Findings

The results of the tests are presented in Table 9-1.

Table 9-1 Summary of Findings

Targeted Attack	Applicability	Considerations	Solution
MAC flooding	No	Macof uses random MAC addresses as source and destination	AP discards frames from a source MAC not in the association table
DHCP starvation	Yes on H-REAP Controller discards bad DHCP requests	The requesting MAC is carried in the DHCP payload	None–rate limiting
Rogue DHCP server	Yes on H-REAP Controller blocks DHCP offers from the WLAN	It is assumed the rogue DHCP server is wireless	None
MIM between wireless clients	Yes on H-REAP Controller blocks GARPs	Traffic does not go through the switch in this case	None
MIM between wireless clients on different APs	Yes on H-REAP Controller blocks GARPs	The hacker can intercept traffic only toward the wire.	DAI with violation
MIM between wireless and wired clients	Yes on H-REAP Not a supported controller configuration	The hacker can intercept traffic only toward the wire.	DAI with violation
IP spoofing	Yes on H-REAP Controller checks IP address and MAC address binding	Encryption over the air is required to prevent identity spoofing	IP Source Guard

Note that Cisco tested only those attacks that are targeted by the CISF features on wired access, and it was always assumed that the attacker was wireless, while the target could be either wired or wireless depending on the scenario considered. Finally, the solution reported in Table 9-1 represents what is currently available using the CISF features on the access switch; when those features do not help, Cisco proposes an alternative solution using features available directly on the access point.

Conclusion

CISF, or at least some of its features, should ideally be implemented on the edge device that provides access to the endpoints. In the case of wireless access, this device is the AP and not the access switch connected to the AP.

As proven by the lab tests, enabling CISF on the access switch still helps prevent or mitigate many of the attacks for which these features were designed. However, the major limitation of this implementation is the failure to stop attacks in which the traffic does not traverse the switch. If the traffic does not pass through the switch, the CISF feature is not activated and does not protect the target.

CHAPTER 10

Cisco Wireless Mesh Networking

The Cisco Wireless Mesh Network solution enables cost-effective, scalable deployment of secure outdoor wireless LANs, providing access to fixed and mobile applications to enhance public safety, efficiency, productivity, and responsiveness.

The design and deployment of Cisco Wireless Mesh Networks solution is too large a topic to be included in this design guide, which is focused on indoor enterprise WLAN deployments. This chapter provides an overview of the Cisco Wireless Mesh Network solution, because its integration into the Cisco Unified Wireless Network Architecture makes it a simple choice for those customers wishing to extend their enterprise WLAN outdoors.

Overview

In the wireless mesh deployment, multiple Cisco 1500 Mesh APs are deployed as part of the same network (see Figure 10-1).

Figure 10-1 Wireless Mesh Deployment

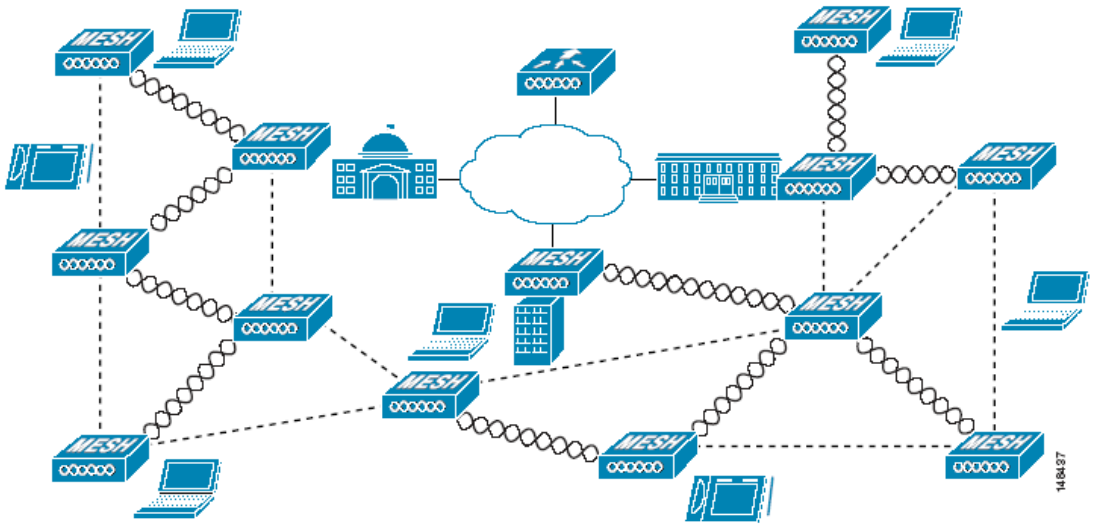

One or more of the Cisco 1500 Mesh APs have a wired connection to their wireless LAN controller, and these are designated as rooftop mesh APs (RAPs). Other Cisco 1500 Mesh APs that relay their wireless connections to connect to the controller are called mesh access points (MAPs). The MAPs use the AWP protocol to determine the best path to their controller through other Cisco 1500 Mesh APs. The various possible paths between the MAPs and RAPs form the wireless mesh that is used to carry traffic from WLAN clients connected to MAPs in that mesh, and also to carry traffic from devices connected to the MAP Ethernet ports.

The WLAN mesh can simultaneously carry two different traffic types: WLAN client traffic and MAP bridge traffic. WLAN client traffic terminates on the WLAN controller, and the bridge traffic terminates on the Ethernet ports of the Cisco 1500 Mesh APs. Mesh membership in the WLAN mesh is controlled in a variety of ways. MAC authentication of the Cisco 1500 Mesh APs can be enabled to ensure that the APs are included in a database of APs that are authorized to use the WLAN controller. Cisco 1500 Mesh APs are configured with a shared secret for secure AP-to-AP intercommunication, and a bridge group name can be used to control mesh membership, or segmentation. The configuration of these features is covered later in this document.

Wireless Backhaul

Cisco 1500 Mesh APs can provide a simple wireless backhaul solution, where the Cisco 1500 Mesh AP is used to provide 802.11b/g services to WLAN and wired clients. This configuration is basically a wireless mesh with one MAP. Figure 10-2 provides an example of this deployment type.

Figure 10-2 Wireless Backhaul

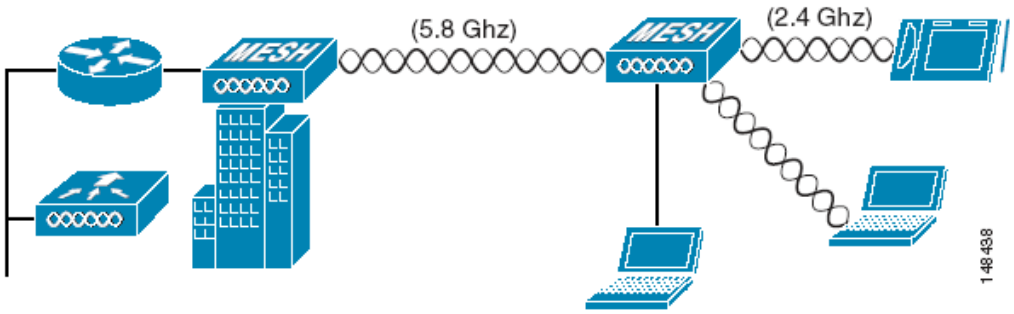

Point-to-Multipoint Wireless Bridging

In the point-to-multipoint bridging scenario, a RAP serving as a root bridge connects multiple MAPs to their associated wired LANs as non-root bridges. By default, this feature is disabled for all MAPs.

If Ethernet bridging is used, you must enable it on the controller for each MAP. Figure 10-3 shows a simple deployment with one RAP and two MAPs, but this configuration is fundamentally a wireless mesh with no WLAN clients. Client access can still be provided with Ethernet bridging enabled, although if bridging between buildings, MAP coverage from a high rooftop might not be suitable for client access.

Figure 10-3 Point-to-Multipoint Wireless Bridging

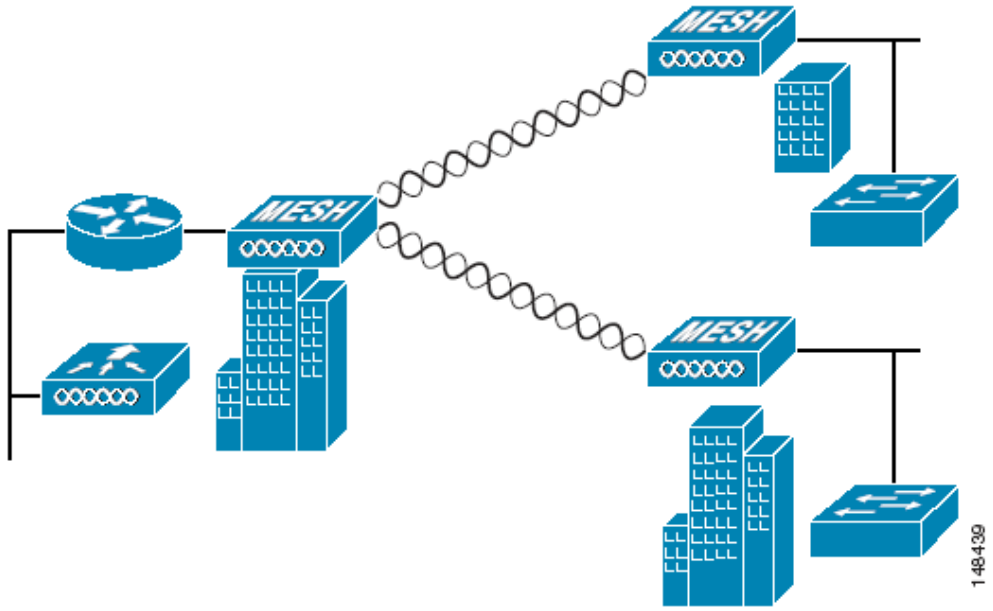

Point-to-Point Wireless Bridging

In a point-to-point bridging scenario, a Cisco 1500 Mesh AP can be used to extend a Layer 2 network by using the backhaul radio to bridge two segments of a switched network, as shown in Figure 10-4. This is fundamentally a wireless mesh with one MAP and no WLAN clients. Just as in point-to-multipoint networks, client access can still be provided with Ethernet bridging enabled, although if bridging between buildings, MAP coverage from a high rooftop might not be suitable for client access.

Figure 10-4 Point-to-Point Wireless Bridging

Wireless Mesh Bridge Connections

The Ethernet ports of the Cisco 1500 Mesh AP are bridged with the wireless mesh acting as a transparent bridge between all Ethernet ports of nodes within that mesh. For example, the simple mesh shown in Figure 10-5 results in a logical multi-port bridge of all Ethernet ports, as shown in Figure 10-6.

Figure 10-5 Simple Mesh Example

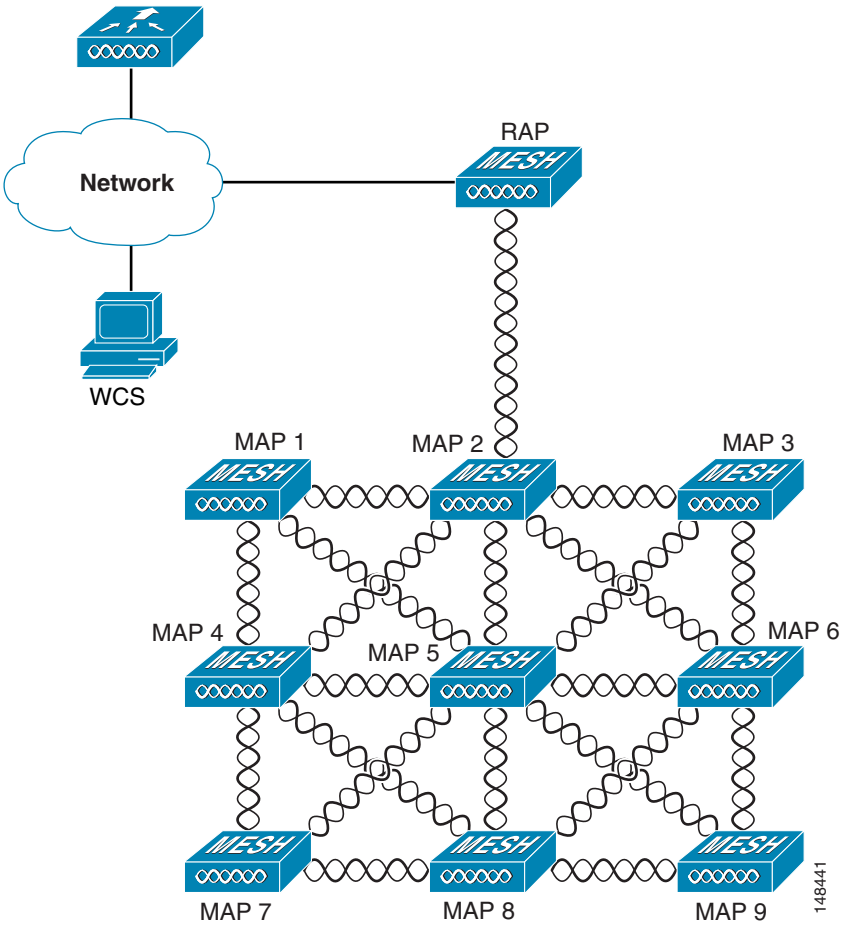

Figure 10-6 Wireless Mesh Virtual Multi-Port Bridge

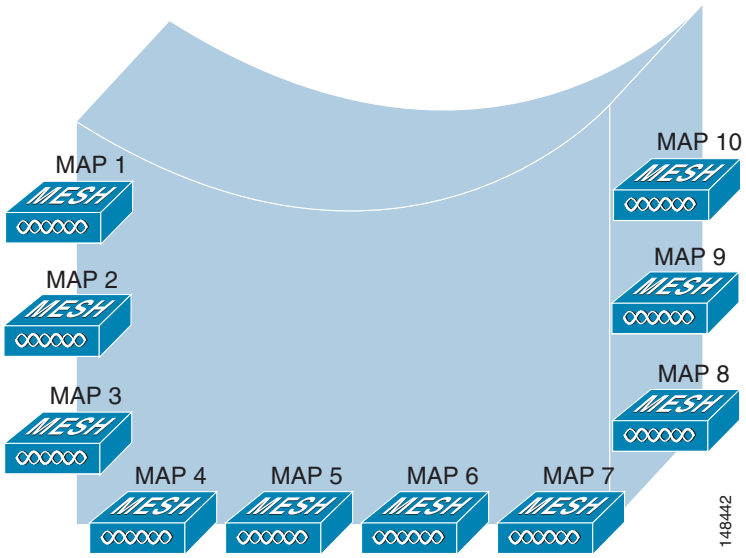

Note that the controller does not participate in this bridging, and that the traffic terminates at the Cisco 1500 AP Ethernet port. Take care to block unnecessary multicast traffic in bridge deployments to prevent wireless backhaul capacity from being consumed unnecessarily.

Also note that for bridged traffic, the controller does not act as a central coordination point. The data traffic for the multipoint bridge is simply bridging traffic through the shortest path calculated by the AWP protocol. The bridge network is transparent to dot1q and Spanning Tree protocols.

Bridge Authentication

When a Cisco 1500 Mesh AP comes up in a mesh, it uses its primary key to authenticate to a parent or a neighboring Cisco 1500 Mesh AP (see Mesh Neighbors, Parents, and Children, page 10-8 for more information). Using this primary key, there is a four-way handshake to establish an AES-encrypted session. The new AP establishes an LWAPP tunnel to the controller and is then authenticated against the MAC filter list of the controller. Next, the controller pushes the bridge shared secret key to the AP via LWAPP, after which it re-establishes the AES-encrypted session with the parent AP.

Wireless Mesh Encryption

As previously described, the wireless mesh bridges traffic between the MAPs and the RAPs. This traffic can be from wired devices being bridged by the wireless mesh, or LWAPP traffic from the mesh APs. This means that the wireless mesh can be carrying traffic that is either clear text or encrypted; this traffic is always AES encrypted when it crosses a wireless link.

The AES encryption is established as part of the mesh AP establishing neighbor relationships with other mesh APs. The bridge shared secret is used to establish unique encryption keys between mesh neighbors.

Simple Mesh Deployment

The key components of the simple mesh deployment design (see Figure 10-6) are as follows:

- A WCS-Key component for the management, operation, and optimization of the mesh network
- An LWAPP controller, which controls the authentication and management of the Cisco 1500 Mesh AP and client WLANs
- A router between the network and the mesh, which provides a Layer 3 boundary where security and policy enforcement can be applied

  The router also provides Layer 2 isolation of the RAP. This is necessary because the RAP bridges traffic from its local Ethernet port to the mesh, so this traffic must be limited to that necessary to support the solution so that resources are not consumed by the unnecessary flooding of traffic.
- A RAP, which provides the path home for the MAP traffic
- Any number of MAPs

Note The RAP wireless connection is toward the center of the MAP mesh, which is an optimal configuration that minimizes the average number of hops in the mesh. A RAP connection to the edge of a mesh would result in an increase of hops.

Figure 10-7 shows one possible logical view of the physical configuration shown in Figure 10-5, with MAP5 as the path home for all other MAPs.

Figure 10-7 Logical View

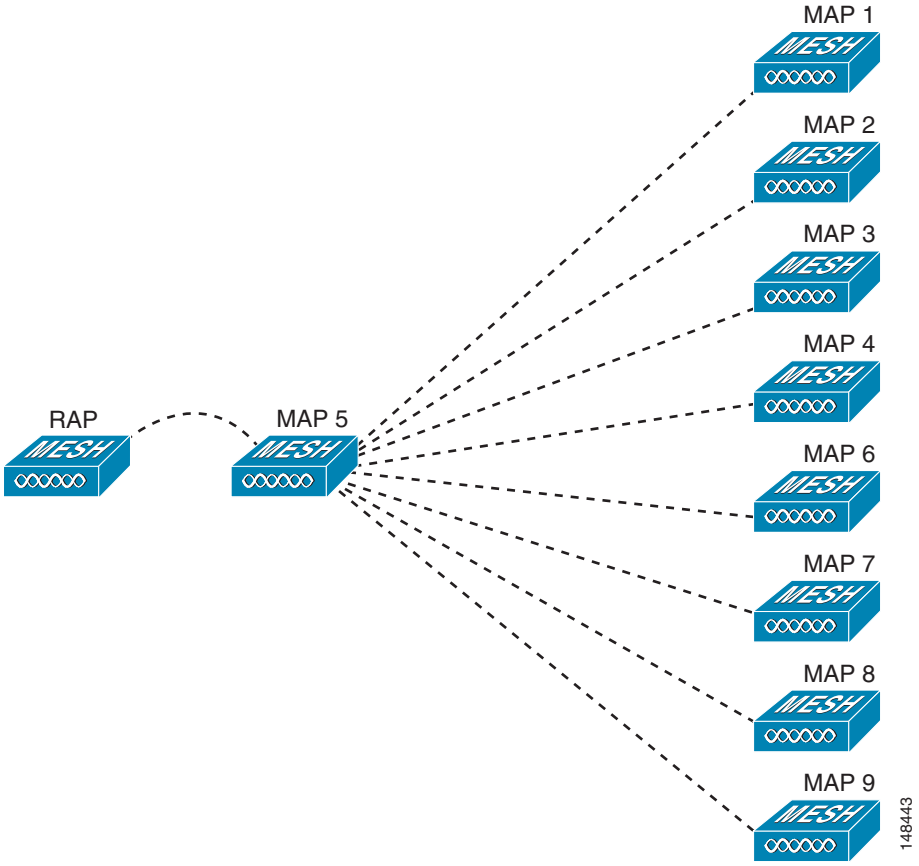

Figure 10-8 shows an alternate logical view, in which the signal-to-noise ratio (SNR) on the diagonal paths to MAP5 is small enough for the MAPs to consider taking an extra hop to get to MAP5.

In both of these cases, MAP5 is the path home for all traffic. Ideally, the coverage from the RAP should be such that other MAPs, such as MAP2 for example, have a path back to the RAP so traffic can be routed via MAP 2 in case of a loss of signal to MAP 5.

Figure 10-8 Unequal Paths

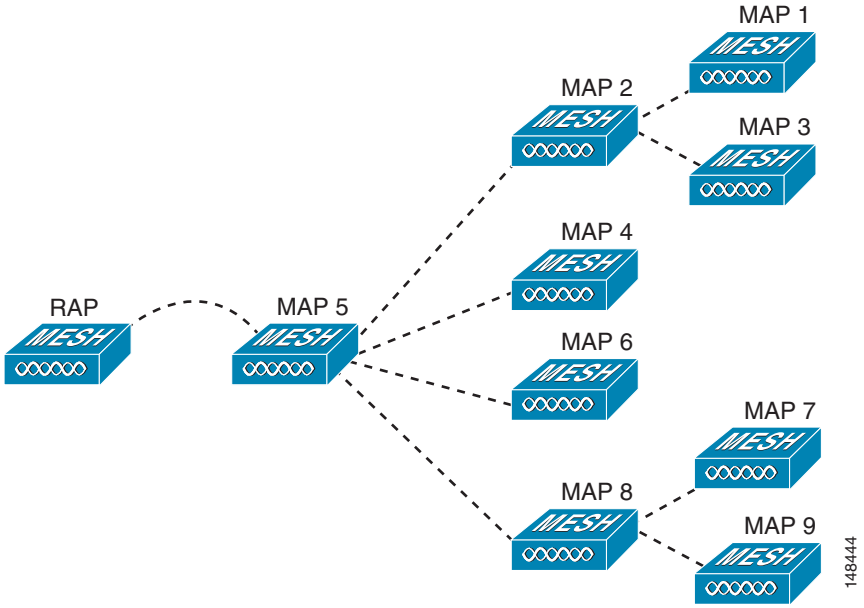

Mesh Neighbors, Parents, and Children

Ease is calculated using the SNR and hop value of each neighbor, and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities. A neighbor within a mesh is an AP that is within RF range that has not been selected as a parent or a child because its “ease” values are lower than another neighboring AP.

A parent AP is one that is selected as the best route back to the RAP based on the best ease values. A parent can be either the RAP itself or another MAP. A child of an AP is an AP that has selected the parent AP as the best route back to the RAP (see Figure 10-9.)

Figure 10-9 Parent, Child, and Neighbor

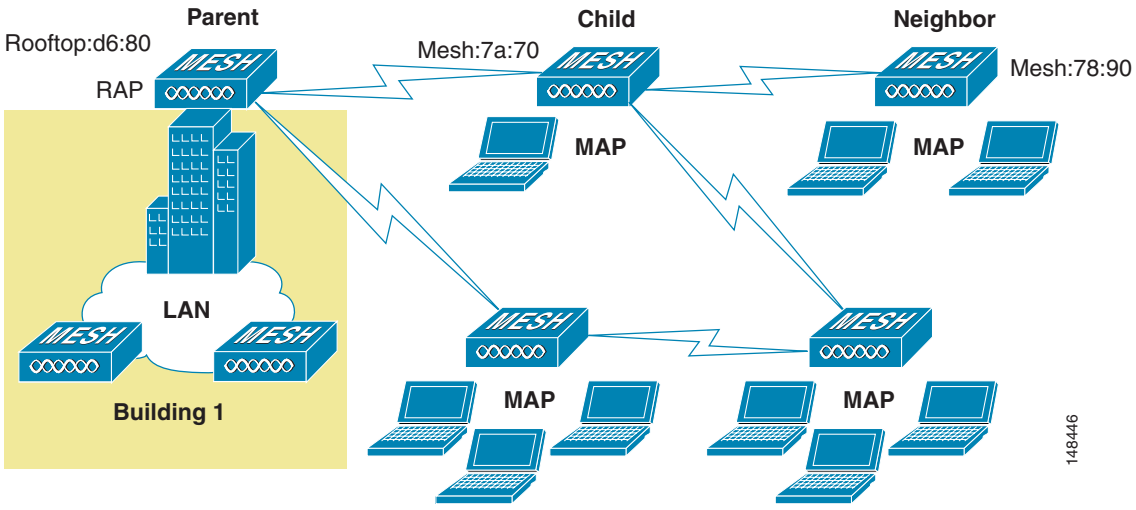

Design Details

Each outdoor wireless mesh deployment is unique, and each environment has its own challenges with available locations, obstructions, and network infrastructure availability, in addition to the design requirements based on users, traffic, and availability. This section covers important design considerations and provides an example of a wireless mesh design.

Wireless Mesh Constraints

When designing and building a wireless mesh network with the Cisco 1500 Mesh AP, there are a number of system characteristics to consider. Some of these apply to the backhaul network design and others to the LWAPP controller design. The recommended backhaul is 18 Mbps. 18 Mbps was chosen as the optimal backhaul rate because it aligns with the maximum coverage of the WLAN portion of the client WLAN of the MAP; that is, the distance between MAPs using 18 Mbps backhaul should allow for seamless WLAN client coverage between the MAPs.

A lower bit rate can allow a greater distance between Cisco 1500 Mesh APs, but there are likely to be gaps in the WLAN client coverage, and the capacity of the backhaul network is reduced. An increased bit rate for the backhaul network either requires more Cisco 1500 Mesh APs, or results in a reduced SNR between mesh APs, limiting mesh reliability and interconnection. The wireless mesh backhaul bit rate, like the mesh channel, is set by the RAP.

The number of backhaul hops is limited to eight, but Cisco recommends that you limit the number of hops to three or four, primarily to maintain sufficient backhaul throughput because each mesh AP uses the same radio for transmission and reception of backhaul traffic. This means that throughput is approximately halved over every hop. For example, the maximum throughput for an 18 Mbps is approximately 10 Mbps for the first hop, 5 Mbps for the second hop, and 2.5 Mbps for the third hop.

There is no current software limitation of how many MAPs per RAP you can configure. However, Cisco recommends that you limit this to 20 MAPs per RAP.

The number of APs per controller is determined by the controller capacity:

- The Cisco 2000 Series Wireless LAN Controller supports up to six APs.

- The Cisco 4400 Series Wireless LAN Controller is fully supported with the Cisco 1500 Mesh AP.
- The Cisco 4402 Series Wireless LAN Controller supports up to 50 APs, two Gigabit ports, and one expansion slot.
- The Cisco 4404 Series Wireless LAN Controller supports up to 100 APs, four Gigabit ports, and two expansion slots.

The number of controllers per mobility group is limited to 24.

Client WLAN

The mesh AP client WLAN delivers all the WLAN features derived by a standard LWAPP deployment for b/g clients with the full range of security and radio management features.

The goals of the client WLAN must be considered in the overall mesh deployment:

- What are the required bit rates?

  Higher bit rates reduce coverage and are limited by the mesh backhaul.
- What throughput is required?
- What are the application throughput requirements, and how many simultaneous clients are expected on a Cisco 1500 Mesh AP?
- What coverage is required?

  Is the coverage between different Cisco 1500 Mesh APs required to be contiguous, or is the mesh deployment a collection of separate active zones?

Design Example

This section provides a sample design of providing WLAN coverage in an urban or suburban area.

Cell Planning and Distance

The starting point is the RAP-to-MAP ratio. There is currently no hard limitation of MAPs per RAPs, but the current recommended maximum number is 20 MAPs per RAP.

For the backhaul, there is a typical cell size radius of 1000 feet. One square mile in feet is 5280^2 square feet, so the number of cells comes out to be nine, and you can cover one square mile with approximately three or four hops (see Figure 10-10).

Figure 10-10 1000 Feet

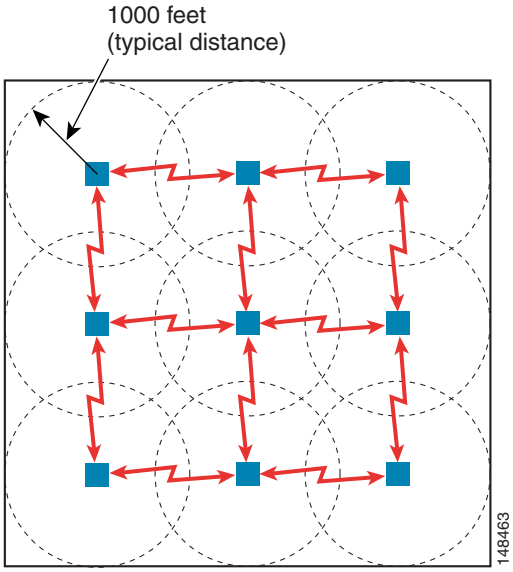

One square mile, 9 cells

For 2.4 GHz, the local access cell size radius is 600 feet. One cell size comes out to be 1.310 x 10 ^6 square feet, so the number of cells is 25 per square mile (see Figure 10-11).

Figure 10-11 600 Feet

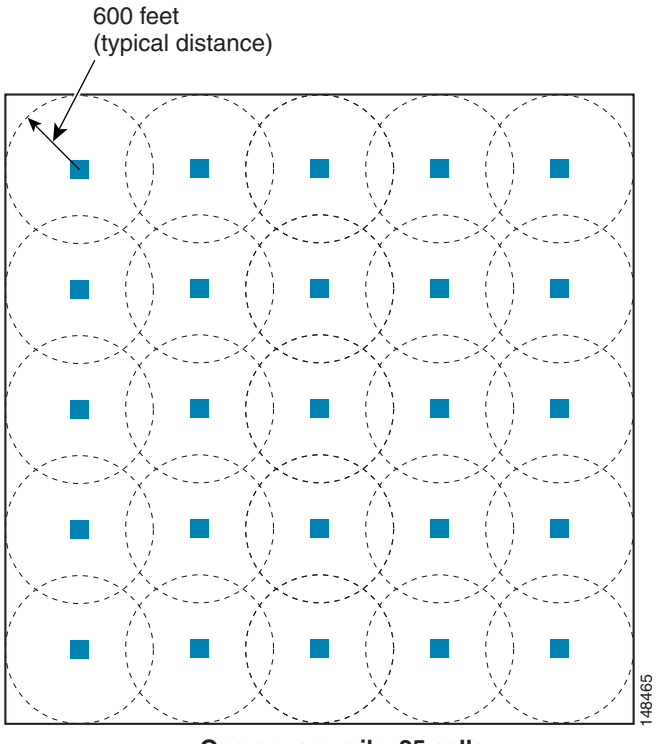

One square mile, 25 cells

The RAP shown in Figure 10-12 is simply a place holder. The goal is to use the RAP location in combination with RF antenna design to ensure that there is a good RF link to the MAPs within the core of the cell. This means that the physical location of the RAP can be on the edge of the cell, and a directional antenna is used to establish a link into the center of the cell.

Figure 10-12 Schematic of the Wireless Mesh Layout

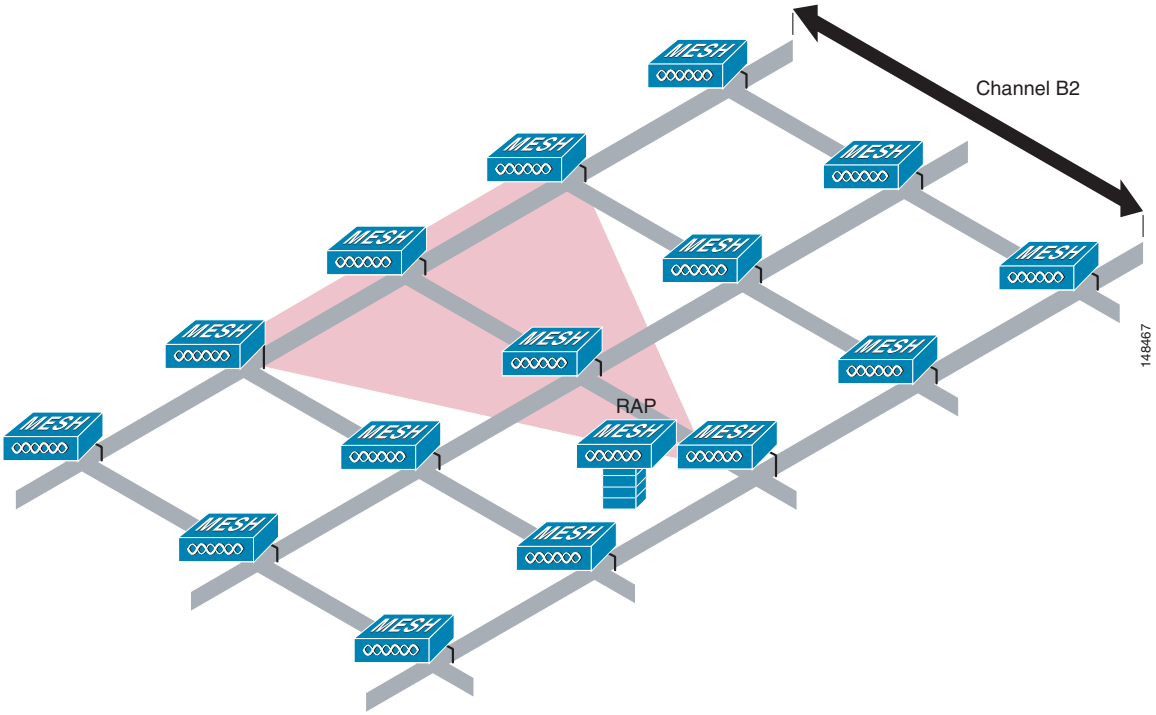

When laying out multiple cells, use channel planning similar to standard WLAN planning to avoid overlapping channels, as shown in Figure 10-13. If possible, the channel planning should also minimize channel overlap in cases where the mesh has expanded to cover the loss of a RAP connection, as shown in Figure 10-14.

Figure 10-13 Laying out Various Cells

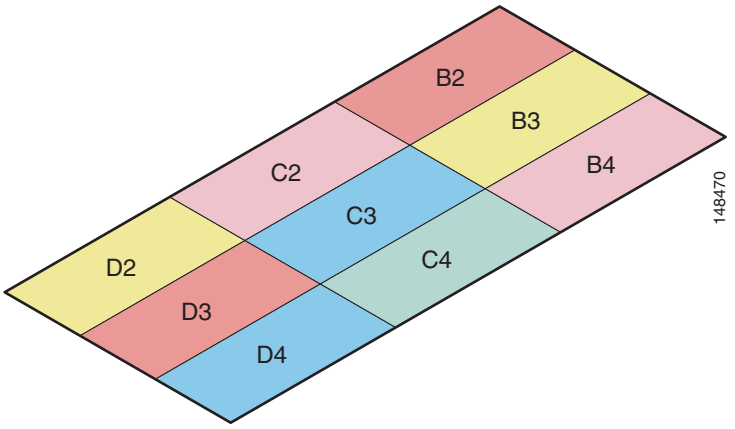

Figure 10-14 Failover Coverage

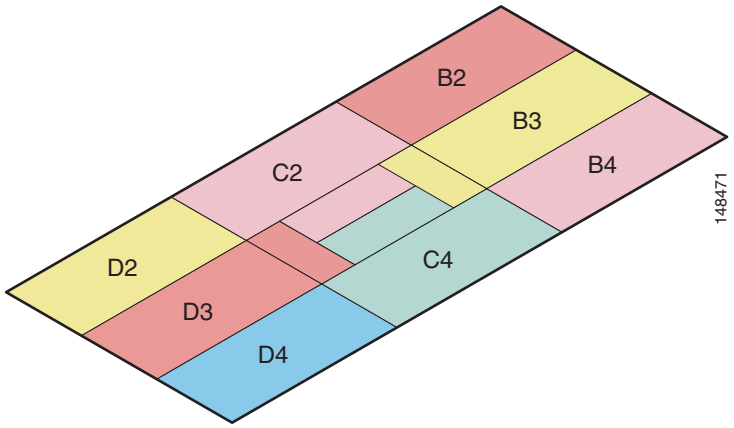

Note Note FCS limitations: with the current release, there is a hard-coded bridging distance limitation of 12000 (2.25 miles) between the Cisco 1500 Mesh APs, even though the radio has a capability to go much further in distance. This distance limitation will be removed in future software releases.

For more information on cell and channel planning including multiple RAP and channels, see the *Cisco 1500 Outdoor Mesh AP Design Guide* at the following URL: http://www.cisco.com/en/US/products/ps6548/products_technical_reference_book09186a008062b50e.html

Controller Planning

A mobility group can have up to 24 WLCs, and WLCs can currently support up to 150 APs; with 24 controllers, this provides a maximum of 3600 APs.

In most cases, the full controller capacity is not normally used in this manner, because some of the controllers are used to increase availability; for example, an n+1 system with 23 active controllers and one controller providing backup services.

Another factor that affects the total number of APs is the wired network connecting the RAPs and controllers. If this network allows the controllers to be equally available to all APs without any impact on WLAN performance, the APs can be evenly distributed across all controllers for maximum efficiency.

If this is not the case, and controllers are grouped into various clusters or PoPs, the overall number of APs and therefore coverage are reduced.

Multiple Wireless Mesh Mobility Groups

Keep in mind that wireless mesh built by the maximum number of controllers in a mobility group is not truly the maximum size of WLAN coverage because this is simply the maximum size of the mobility group. The WLANs that are part of a mobility group can be replicated in another mobility group, and a WLAN client is able to roam between these mobility groups.

When roaming between mobility groups, the roaming can be Layer 2 roaming or Layer 3 roaming, depending on the network topology behind the wireless mesh networks.

Increasing Mesh Availability

In the previous section, a wireless mesh cell of one square mile was created and then built upon. This wireless mesh cell has similar properties to the cells used to create a cellular phone network; that is, although the technology might define the maximum size of the cell, smaller cells can be created to cover the same physical area, providing greater availability or capacity. This is done by adding RAPs to the cell. Just as in the larger mesh deployment, the decision is whether to use RAPs on the same channel, as shown in Figure 10-15, or to use different channels, as shown in Figure 10-16. The addition of RAPs into an area adds capacity and resilience to that area.

Figure 10-15 Two RAPs per Cell with the Same Channel

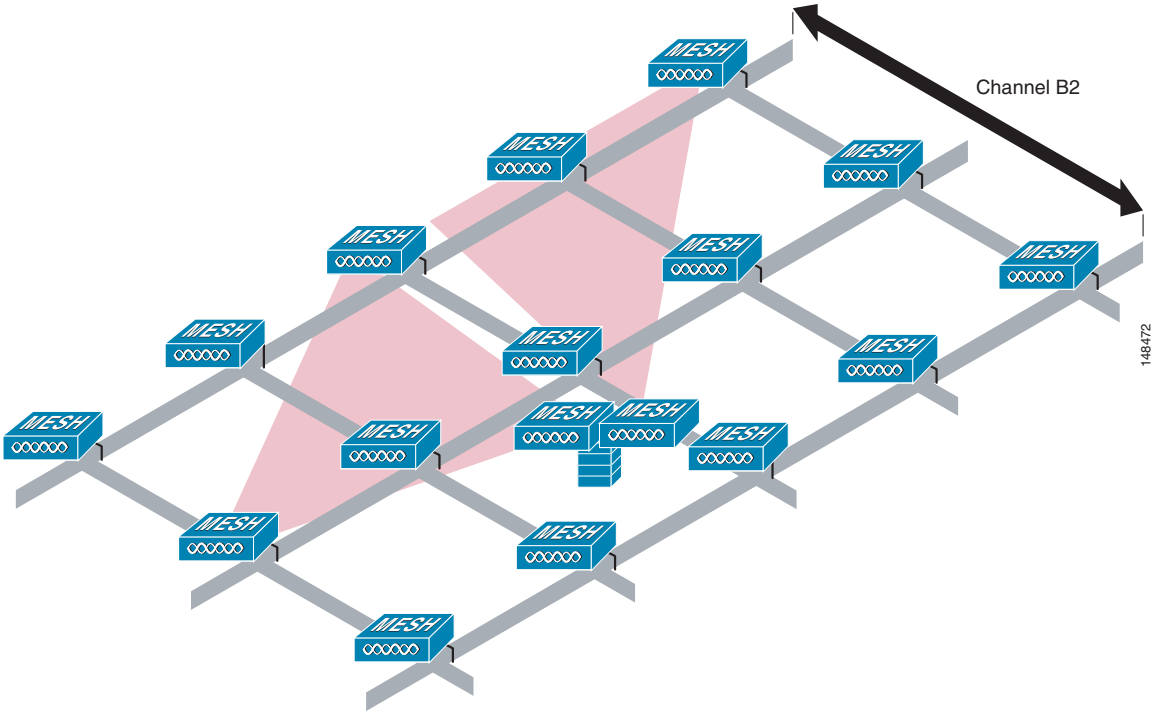

Figure 10-16 Two RAPs per Cell on Different Channels

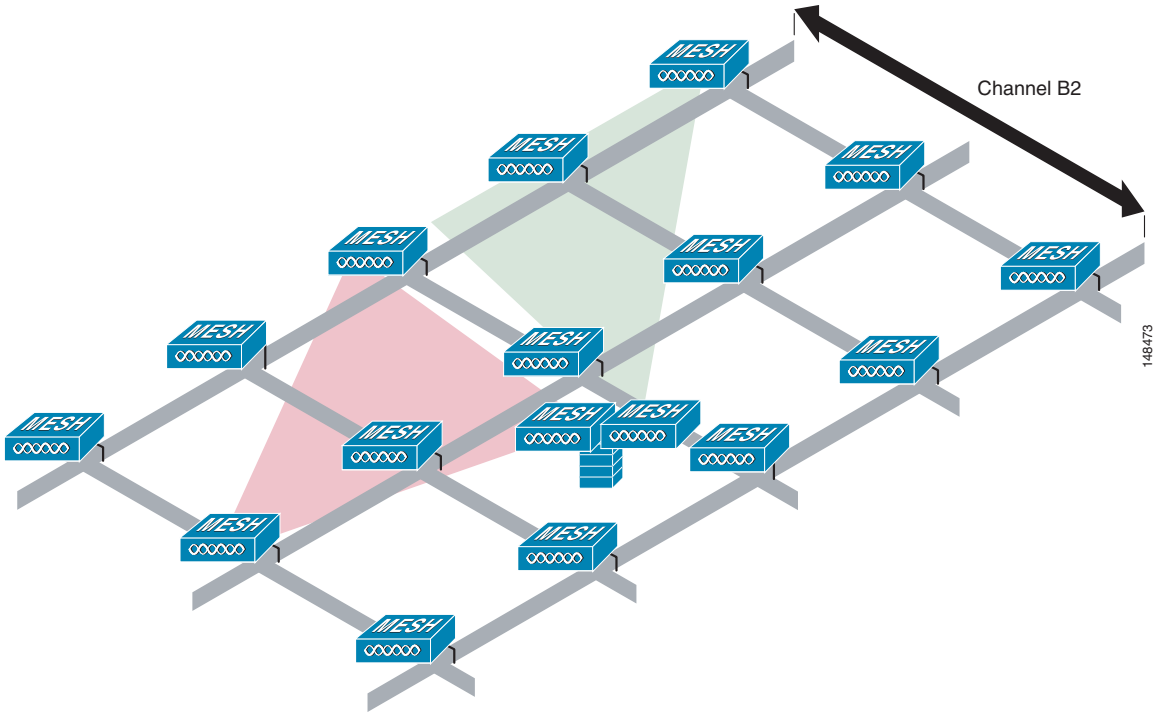

Layer 2 Versus Layer 3 Encapsulation

Cisco generally recommends using Layer 3 encapsulation because it gives greater flexibility in RAP and controller placement. Even if it is possible to put the RAP and its associated controllers on the same subnet, Cisco recommends that the RAP and the controllers be separated by a router hop, because this controls the Layer 2 traffic going into the RAP Ethernet interface, and simplifies the network design if more RAPs or controllers need to be added.

Multiple RAPs

If multiple RAPs are to be deployed, the purpose for deploying these RAPs needs to be considered. If the RAPs are being deployed to provide hardware diversity, the additional RAPs should be deployed on the same channel as the primary RAP to minimize the convergence time in a scenario where the mesh transfers from one RAP to another. When planning RAP hardware diversity, remember the 32 MAPs per RAP limitation.

If the additional RAPs are being deployed to primarily provide additional capacity, deploy the additional RAPs on a different channel to its neighboring RAPs to minimize the interference on the backhaul channels.

If the mesh AP bridging features are being used with multiple RAPs, these RAPs should all be on the same subnet to ensure that a consistent subnet is provided for bridge clients.

If you build your mesh with multiple RAPs on different subnets, MAP convergence times increase if a MAP has to failover to another RAP on a different subnet. One way to limit this from happening is to use different BGNs for segments in your network that are separated by subnet boundaries.

Multiple Controllers

The consideration in distance of the LWAPP controllers from other LWAPP controllers in the mobility group, and the distance of the LWAPP controllers from the RAPs, is similar to the consideration of an LWAPP WLAN deployment in an enterprise.

There are operational advantages to centralizing LWAPP controllers, and these advantages need to be traded off against the speed and capacity of the links to the LWAPP APs and the traffic profile of the WLAN clients using these APs.

If the WLAN client traffic is expected to be focused on particular sites such as the Internet or a data center, centralizing the controllers at the same sites as these traffic focal points gives the operational advantages without sacrificing traffic efficiency.

If the WLAN client traffic is more peer-to-peer, a distribute controller model might be a better fit; it is likely that a majority of the WLAN traffic are clients in the area, with a smaller amount of traffic going to other locations. Given that many peer-to-peer applications can be sensitive to delay and packet loss, it is best to ensure that traffic between peers takes the most efficient path.

Given that most deployments see a mix of client server traffic and peer-to peer traffic, it is likely that a hybrid model of LWAPP controller placement is used, where points of presence (PoPs) are created with clusters of controllers placed in strategic locations in the network.

In all cases, remember that the LWAPP model used in the wireless mesh network is designed for campus networks; that is, it expects a high-speed, low-latency network between the LWAPP APs and the LWAPP controller.

Indoor WLAN Network to Outdoor Mesh

Mobility groups can be shared between outdoor mesh networks and indoor WLAN networks. It is also possible for a controller to control indoor LWAPP APs and Cisco 1500 Mesh APs simultaneously. The same WLANS are broadcast out both the indoor AP as well as the Cisco 1500 Mesh APs.

Adding outdoor mesh networks in an enterprise campus network enables new mobility applications. Voice clients are able to take their calls in motion from one building to another on the same WLANs they have inside the buildings. Security can also benefit from this new outdoor wireless network. Video streaming to security officer vehicles from outdoor and indoor cameras, access to data from the network, as well as transmitting data voice and video from the vehicle are all enabled. Because of the fast secure roaming enabled by the controllers and APs, outdoor applications can now move around the campus.

Outdoor Mesh Controllers

You can either associate your outdoor mesh APs with the same controllers you are using for your indoor enterprise network, or you can have a separate controller or pool of controllers just for your outdoor network. If you choose to use the same controllers for both your indoor and outdoor network, take into consideration that there might be some WLANs you do not want on your outdoor network. In this case, use the WLAN override option under the b radio configuration on your outdoor mesh APs to disable specific WLANs. For example, you might not want to have your guest access network enabled on the outdoor network. Figure 10-17 displays this example.

The other option is to have a separate pool of controllers just for your outdoor network. This way allows you to configure and control all your WLANs on a per controller basis without having to disable certain WLANs on every outdoor MESH AP. If you take this approach, you must configure the mobility groups between your designated outdoor controllers and your indoor controllers to provide seamless roaming for clients between the two networks.

Figure 10-17 Indoor Outdoor WLANs

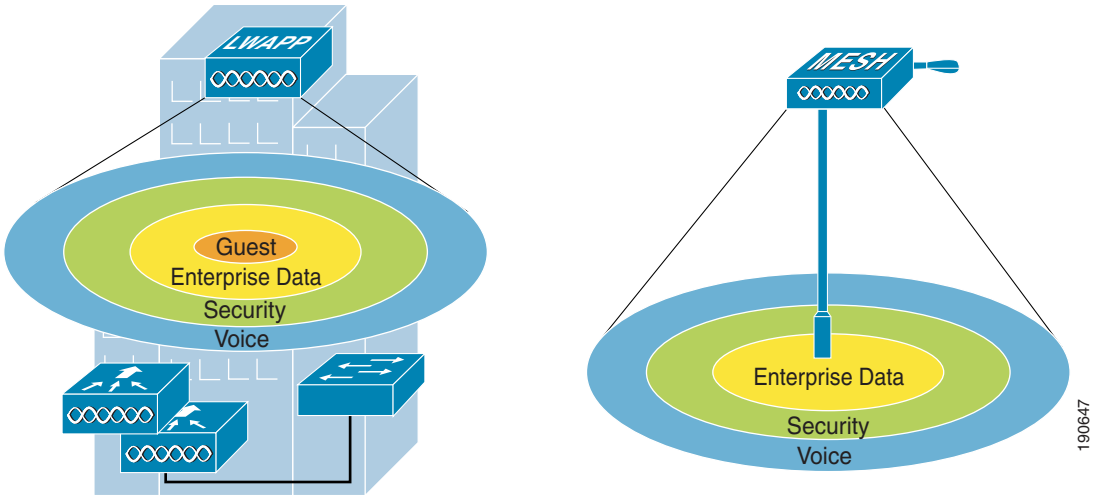

Connecting the Cisco 1500 Mesh AP to your Network

The wireless mesh has two locations where traffic terminates on the wired network. The first location is where the RAP attaches to the wired network, and where all bridged traffic connects to the wired network. The second location is where the LWAPP controller connects to the wired network; this is where WLAN client traffic from the mesh network connects to the wired network. This is shown schematically in Figure 10-18. The WLAN client traffic from LWAPP is tunneled at Layer 2, and matching WLANs should terminate on the same switch VLAN as where the controllers are co-located. The security and network configuration for each of the WLANs on the mesh depend on the security capabilities of the network to which the controller is connected.

Figure 10-18 Mesh Network Traffic Termination

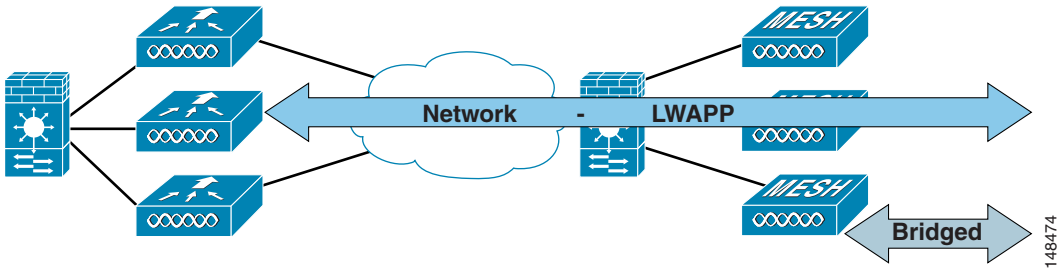

Indoor LWAPP APs do not need a firewalled connection like the outdoor APs. This is because their interfaces are on a secure indoor network and all traffic is tunneled to the controllers, unlike the outdoor APs that can be used for a bridging application.

Physical Placement of Outdoor Mesh APs

When choosing a location for your APs on the campus, keep in mind issues such as building height obstructions, light pole locations, and power options. There are light poles on most enterprise campuses, but not all of them are equipped with a electric eye. Make note of what types of light poles you have and

options for taping power. When placing the roof top AP, a directional antenna might be of use to direct coverage to a specific MAP or group of MAPs designated as the first hops into the mesh. If you plan to use an omni-directional antenna for the RAP, make sure to mount it towards the edge of the building so the radio coverage is not blocked by the edge of the building. Figure 10-19 shows coverage concerns between the RAP and MAPs in the mesh.

Figure 10-19 AP Placement

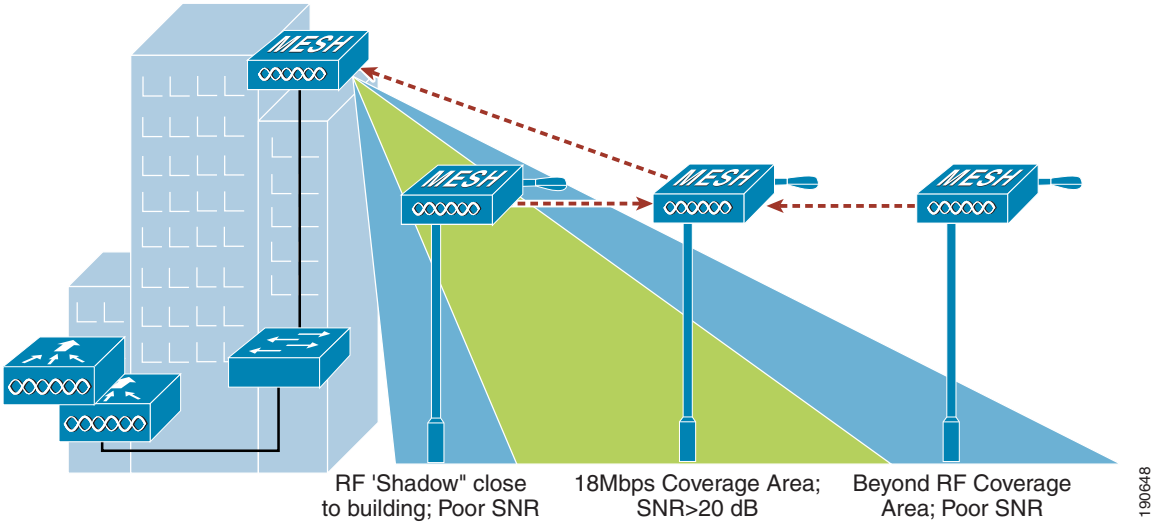

For more information on MESH QoS, Cisco Adaptive Wireless Path (AWP) protocol, mesh traffic flow, Ease, and CLI commands, see the *Cisco 1500 Mesh AP Deployment Guide* at the following URL: http://www.cisco.com/en/US/products/ps6548/products_technical_reference_book09186a008062b50e.html .

VoWLAN Design Recommendations

This chapter provides design considerations when deploying voice over WLAN (VoWLAN) solutions. WLAN configuration specifics may vary depending on the VoWLAN devices being used and the WLAN design. This chapter provides more details about key RF and site survey considerations that are generally applicable to VoWLAN deployments, which were introduced in Chapter 3, "WLAN Radio Frequency Design Considerations."

Antenna Considerations

The more demanding network requirements of VoWAN impacts WLAN planning at all levels, down to the choice of antenna. Key antenna considerations are as follows:

- Access point (AP) antenna selection
- Antenna placement
- Handset antenna characteristics

AP Antenna Selection

Cisco recommends a diversity ceiling-mount antenna for voice applications. Ceiling mounted antennas offer a quick and easy installation. More importantly, they place the radiating portion of the antenna in open space, which allows the most efficient signal propagation and reception. Cisco recommends that all antennas be placed 1 to 2 wavelengths from highly reflective surfaces such as metal. The 2.4 GHz wave is 4.92 inches (12.5 cm) and the 5 GHz is 2.36 inches (6 cm). The wavelength separation between the antenna and reflective surfaces allows the AP radio a better opportunity to receive a transmission, and reduces the creation of nulls when the AP radio transmits. Orthogonal frequency-division multiplexing (OFDM) used by 11g and 11a helps to mitigate problems with reflections, nulls, and multipath; however, good antenna types and placement provide a superior solution. The ceiling tile itself is a good absorber of signals transmitted into the area above the ceiling and reflected back into the coverage area.

Antennas come in many types and enclosures; no single type or module of antenna is best for all applications and locations. For additional information on the performance of various antenna types, and the part numbers of Cisco Aironet antennas, see the Cisco Aironet antenna guide at the following URL: http://www.cisco.com/en/US/products/hw/wireless/ps469/products_data_sheet09186a008008883b.html.

When attaching antennas to an AP, Cisco recommends using the Cisco AIR-ANT5959 for 2.4 GHz and AIR-ANT5145V-5 for 5 GHz for indoor voice applications. These two antennas provide the following advantages:

- Low gain omni-diversity
- Reduced uptilt, which reduces the coverage that may sweep into the floor above, and also reduces the reflections that may come from air ducts and other metal objects above the ceiling tile
- Easy attachment to the T-bar on most ceiling tiles

Higher gain antennas spread the signal on the horizontal plane, which creates a larger cell that also picks up more noise. This results in a smaller signal-to-noise ratio (SNR), which increases the packet error ratio. SNR is defined by the following two criteria:

- Signal—The radiated energy transmitted from one radio that can be received uninterrupted by another radio. For WiFi radios, this means that the transmitting radio is sending 802.11 protocol packets that the receiving radio can decode.
- Noise—Transmitted energy in the frequency range of the receiving radio which cannot be decoded by that radio.

The larger the difference in energy between the protocol packet and the background noise, the better the reception of the protocol packet and the lower the packet error rate and bit error rate. Coverage area design involves using channels to create the lowest possible packet error rate while maintaining a high call capacity.

Higher gain antennas can also reduce the number of calls on a WiFi channel because of the increased coverage area. The Cisco AP1130 Series provides the same type of antenna design concept. For voice, a ceiling-mounted antenna is preferred over a wall-mounted patch because the human head and body attenuate 5 dB of the signal (see Figure 11-1). Most floors attenuate 7 dB of the signal.

Figure 11-1 Head and Hand Attenuation

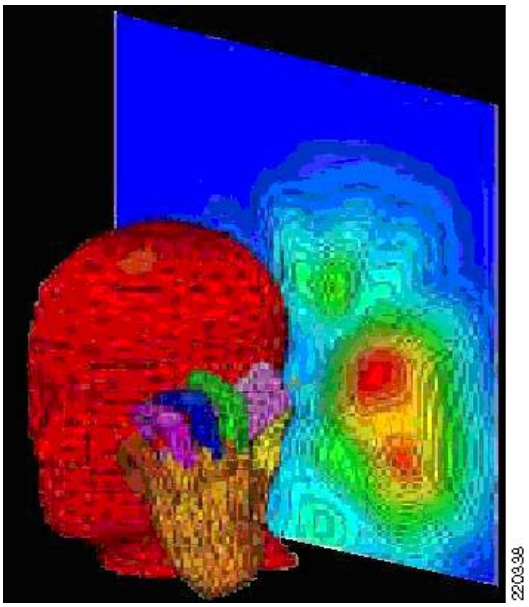

Antenna Positioning

A ceiling down signal has a more direct path to a phone. The recommended coverage cell size takes into consideration the signal loss because of the attenuation of the head and other obstacles. It is important to understand that the gain of antennas is reciprocal: gain applies equally to reception and transmission. Antenna gain is not an increase in transmitted power; the radio produces the transmitted power, and the antenna is a passive device. Gain is derived from focusing the signal of the radio into a direction, and plane and beam width.

For a further discussion of WLAN RF planning, see Chapter 3, "WLAN Radio Frequency Design Considerations."

Handset Antennas

The Cisco Unified Wireless IP Phone 7920 and 7921 have antennas that extend from the main body of the phone. The way they are held in the hand does not influence possible signal attenuation from the hand.

For phones that have the antenna inside the body of the phone, the way the phone is held in the human hand can influence signal attenuation by 4 dB. In some cases, a phone held against the head with the hand covering the antenna can result in a signal drop of 9 dB. The general rule for indoor deployments is that every 9 dB of signal loss reduces the coverage area in half. Figure 11-1 shows an example of the difference in radiating power from a handset when held to the head.

Handsets using the 2.4 GHz spectrum generally do not have diversity antennas for the 2.4 GHz spectrum because the 2.4 GHz wave is nearly five inches, and no combination of diversity antenna options improve signal reception. Therefore, the only improvement in link quality is at the AP. To provide the best quality of link between the phone and the AP, the AP needs to be in its default configuration of diversity-enabled and have diversity antenna support.

Note that 802.11a handsets such as the Cisco 7921 may have a diversity antenna solution for the 11a radio.

Channel Utilization

The 802.11, 802.11b, and 802.11g use the same 2.4 GHz band. The existing WiFi protocols in the 2.4 GHz band need to interoperate with each other, which brings additional overhead, reducing channel throughput. Many sites already have products using the WiFi 2.4 GHz band. In addition, many other products use the same 2.4 GHz frequencies used by WiFi. Other products include Bluetooth, wireless handsets, video game controllers, surveillance cameras, and microwave ovens. Because of the existing use of the channel-limited 2.4 GHz bands, the crowding in the 2.4 GHz spectrum and the constraints in a channel allocation mean that you should consider the 5 GHz WiFi band for new VoWLAN deployments. The channels available in 5 GHz are generally free of use at most sites (see Figure 11-2). Use of the UNII-2 channels for VoWLAN traffic requires the absence of radar. Cisco therefore recommends that there should be extra testing at a new site to see whether any channel in UNII-2 should be blocked out by configuration. The reason for this is if an AP detects radar during normal use, it must leave the channel within ten seconds.

Figure 11-2 Typical Office Channel Utilization for 2.4 GHz and 5 GHz

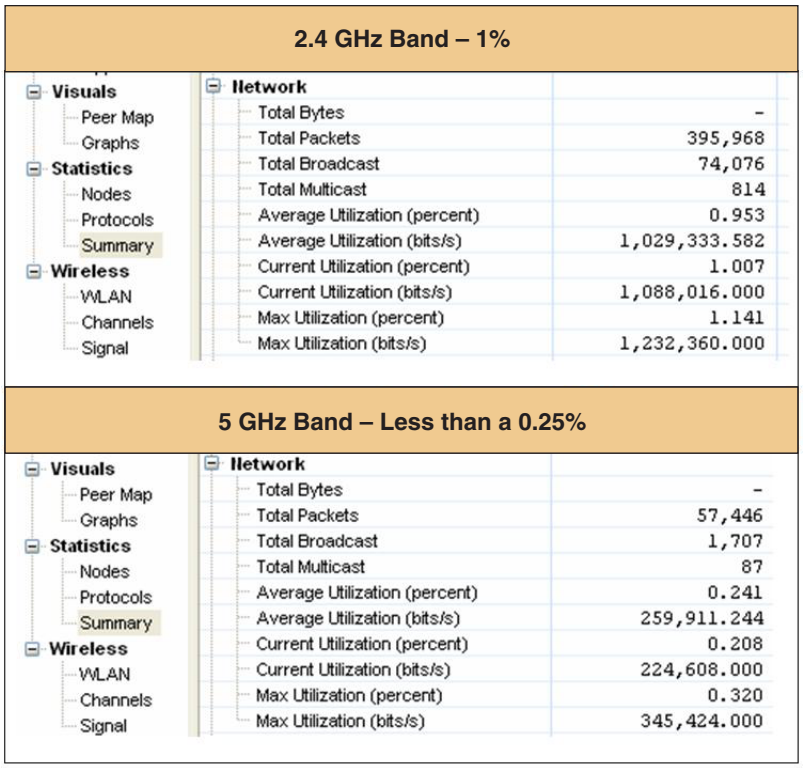

Before the installation of the Cisco Unified Wireless Network, a site can be tested for channel interference and utilization with tools from AirMagnet, Wild Packets, Cognio, and others. The Wireless Control System (WCS) AP On-Demand Statistics Display report provides a spectrum review of the following:

- Noise by channel
- Interference by channel
- Client count versus RSSI
- Client count versus SNR
- Channel radar detection versus time

An example of these statistics can be found in Chapter 8, "Cisco Unified Wireless Control System." Individual client statistics of client RSSI history, client SNR history, and bytes sent and received in Kbps and per second are illustrated in Appendix E, "Sample Monitor > Devices > Access Points Reports."

Dynamic Frequency Selection (DFS) and 802.11h Requirements of the APs

The Federal Communications Commission (FCC) of the United States, the European Telecommunications Standards Institute (ETSI), and other regulatory agencies have requirements regarding the use of radio frequencies. Portions of the 5 GHz band have been and are currently being used for radar, such as weather radar. Although most 5 GHz radar systems generally use high frequencies with shorter wavelengths, there are still systems in place that overlap with some WiFi UNII-2 bands. In 2006, the FCC opened the frequencies in the 5470–5725 MHz to unlicensed use. With these additional frequencies came the requirement of maintaining an interference-free AP configuration. The AP must

constantly monitor for radar pulses (typically from military, satellite, and weather stations), and must automatically switch to a "clean" channel if radar is detected. When radar is detected, the system must do the following:

- Stop packet transmission within 200 ms
- Stop control transmissions within 10 seconds
- Avoid transmission on the channel for 30 minutes
- Scan the new channel for 60 seconds before transmission

Because of the radar requirements in the UNII-2, you should conduct a test for radar before going live with voice applications, because the behavior required when radar is detected may impact voice call quality. WCS reporting of detected radar signals is shown in Chapter 8, "Cisco Unified Wireless Control System.". Cognio Spectrum is also an excellent tool to test for radar. If radar is detected during such a test, the APs can then be configured to not use those channels.

Channels in the 5 GHz Band

Figure 11-3 shows the FCC 802.11a channel assignments. The DFS requirement includes the four original UNNI-2 channels (52–64) and the new eight channels (100–116 and 132–140). The 5 GHz band now has 20 channels. These are non-overlapping channels, which means that they can all be co-located. 2.4 GHz has only three non-overlapping channels. A design allowing co-located channels in a coverage area aggregates the number calls obtainable in a coverage area.

Figure 11-3 802.11a Channel Allocation

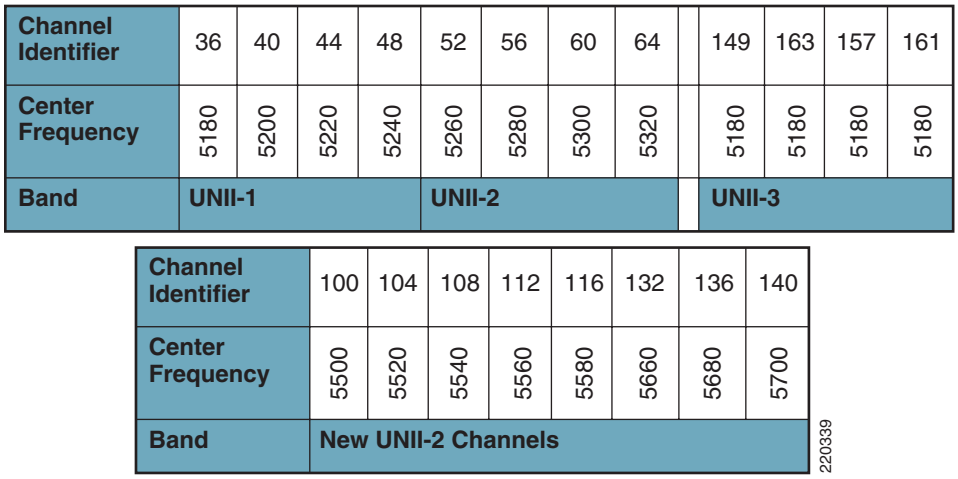

The coverage design based on channels may be done to a single floor, as shown in Figure 11-4. In a multi-floor site, the channels can be separated between floors to reduce the possibility of co-channel interference.

Figure 11-4 Single Floor Channel Design

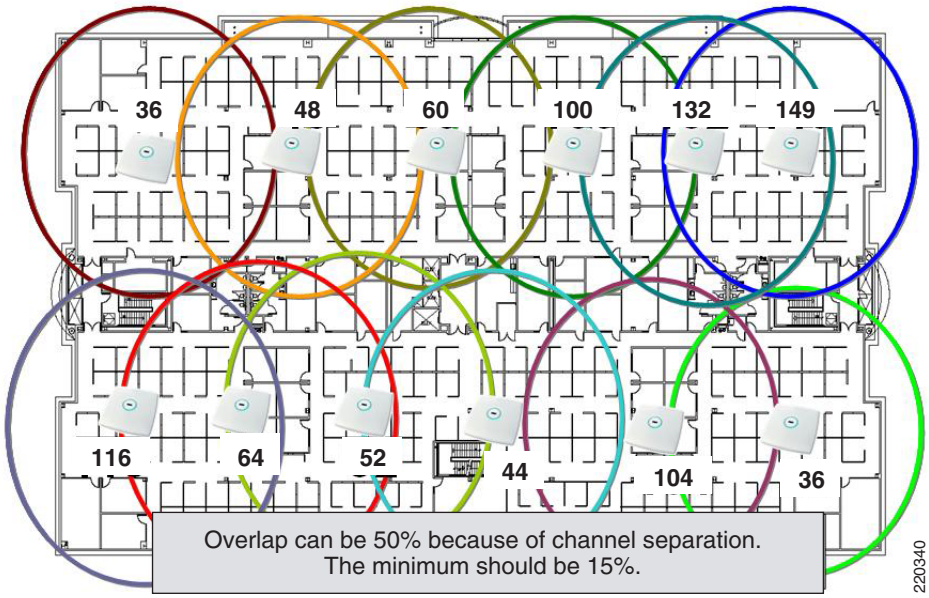

Figure 11-5 illustrates the vertical channel separation.

Figure 11-5 Vertical Channel Separation

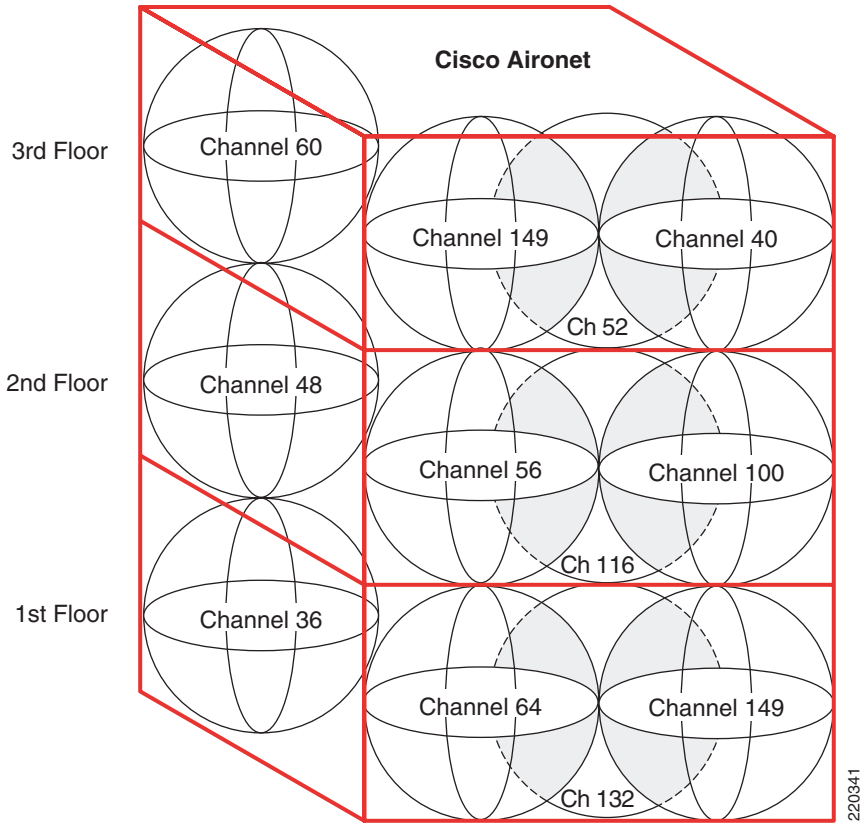

Call Capacity

The number of calls on a WiFi channel is limited by a number of factors. First, the media used by the AP and VoWLAN clients is the RF spectrum. RF spectrum cannot be shielded from interference like shielded twisted-pair CAT 5 cable or fiber. The closest WiFi comes to segmentation is channel separation. This open shared media creates the possibility for high packet loss in 802.11. Most of this packet loss is addressed through retransmission of 802.11 frames, which in turn causes jitter. Figure 11-6 illustrates the packet loss relationship as a mean opinion score (MOS).

Figure 11-6 Effective Packet Loss Graphic

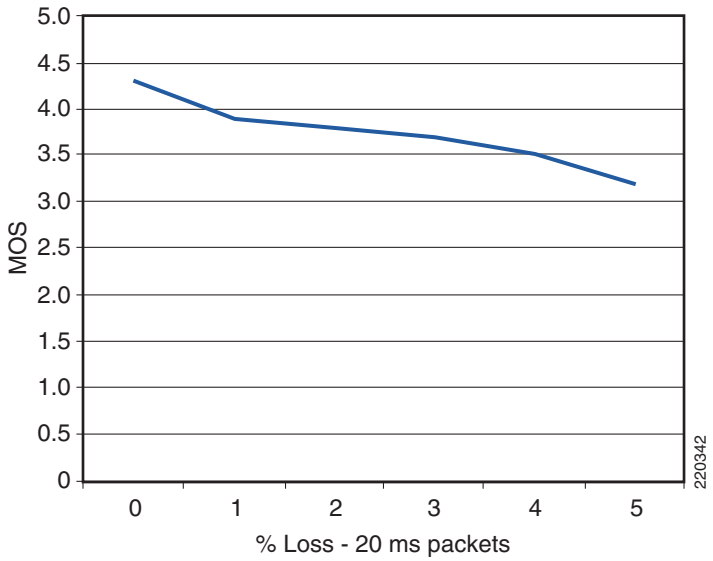

In 802.11a as well as 802.11g, the highest coverage range is achieved by the lowest data rate, which is 6 Mbps. The lowest packet error rate is also at 6 Mbps, for the same given power level. Note that in 802.11b, the lowest data rate is 5.5 Mbps; and in 802.11, the lowest rate is 1 Mbps.

A successful coverage area for voice is an area that maintains a packet error rate of 5 percent or less. The MOS scores are ranked as follows:

- 4.4—Top G.711 MOS score
- 4.3–4.0—“Very satisfied” to “satisfied”
- 4.0–3.6—“Some users satisfied”

Figure 11-6 shows that a packet error rate of 5 percent reduces the MOS to a level of “some users satisfied” quality of speech.

The coverage area edge for a phone is where the coverage area drops to a “very satisfied” MOS. This coverage area edge is referred to as the *cell edge* in this chapter. A cell edge with a 1 percent packet error rate at survey is needed for voice because of the likelihood of multiple phones clients, data clients, co-channel interference, and other un-accounted for interferers. Cell edge and coverage design are defined in detail in other sections of this chapter.

If 802.11 and 802.11b are not required to support legacy 2.4 GHz WiFi clients, Cisco recommends disabling the rates of 1, 2, 5.5, and 11. If those rates are disabled, one or more 802.11g data rates must be set to “required”. The data rate of 6 is generally the recommended data rate to be set to “required”, but this depends on the cell size design requirements, which may require using a higher bit rate. If possible, an 802.11g-only network is recommended rather than a 802.11b/g network. Most data clients

and phone clients recognize the data rates advertised by the AP in its beacons and probe response. Therefore, the clients send their management, control, multicast, and broadcast packets at the "required" data rates as advertised by the AP. The clients can send their unicast packets at any of the data rates advertised by the AP. Generally, those unicast packets are sent at a data rate that provides the highest reliable data rate for the link between the AP and client. The AP is capable of sending unicast packets at a data rate that is unique to each client link.

SNR is an important consideration for packet reception. The receiving radio is either the AP radio or the phone radio. The SNR is not likely to be the same at both radios of the link. SNR and multipath interference must be considered at the AP and at the coverage area edge. Path loss can be assumed to be the same at both ends of the link.

Cisco recommends for voice applications that the cell edge be determined by using the actually phone at the desired data rate. The voice packets sent between the AP and the phone in WiFi applications are generally unicast RTP G711 packets with a typical size of 236 bytes. The Real-Time Transport Protocol (RTP) packet is based on UDP and IP protocols, and therefore RTP is connectionless. The signal strength, SNR, data rate, and error rates of the phone call can be seen from the AP statistics, either on the autonomous AP or the Lightweight Access Point Protocol (LWAPP) controller. A sample of a phone client's cell edge dBm values for 11g and 11a are shown in Figure 11-7 and Figure 11-8. The call stream statistics are shown in Figure 11-9. The stream metrics can be viewed on the WCS after the voice metrics are enabled. The path to enable the metrics is Configure > Controller > ipaddress > 802.11bg > Voice Parameters > Enable Voice Metrics.

Figure 11-7 11g Client Statistics

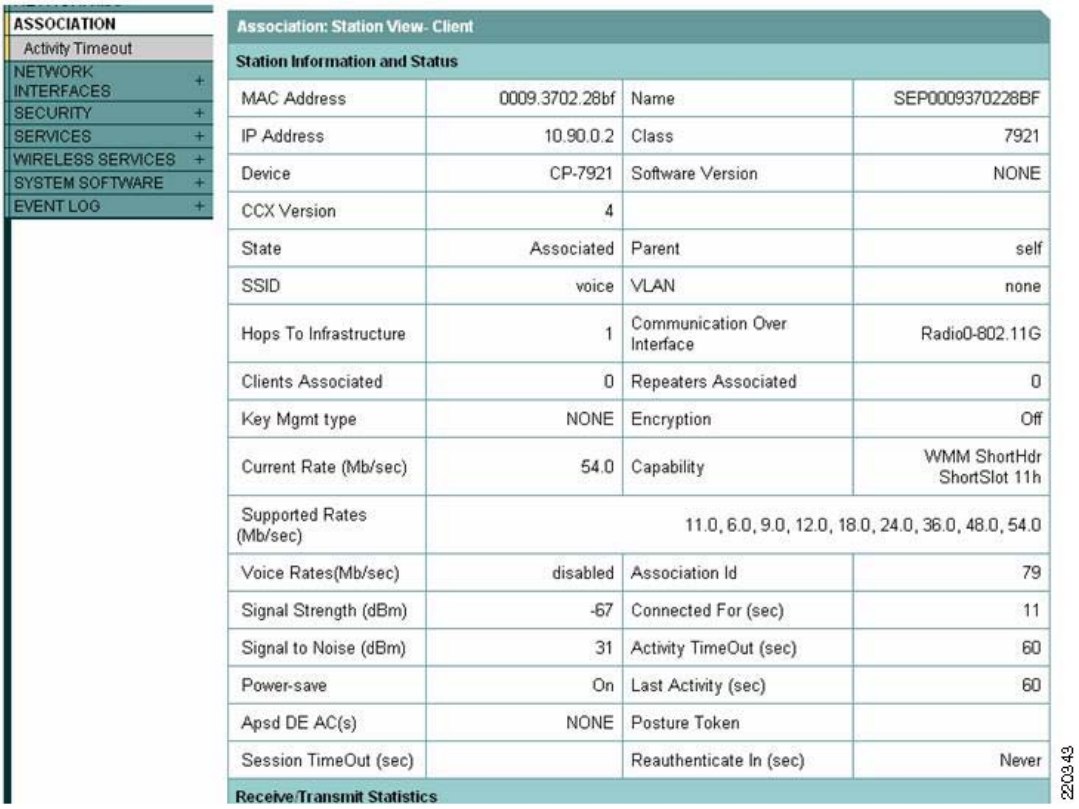

Figure 11-8 11a Client Statistics

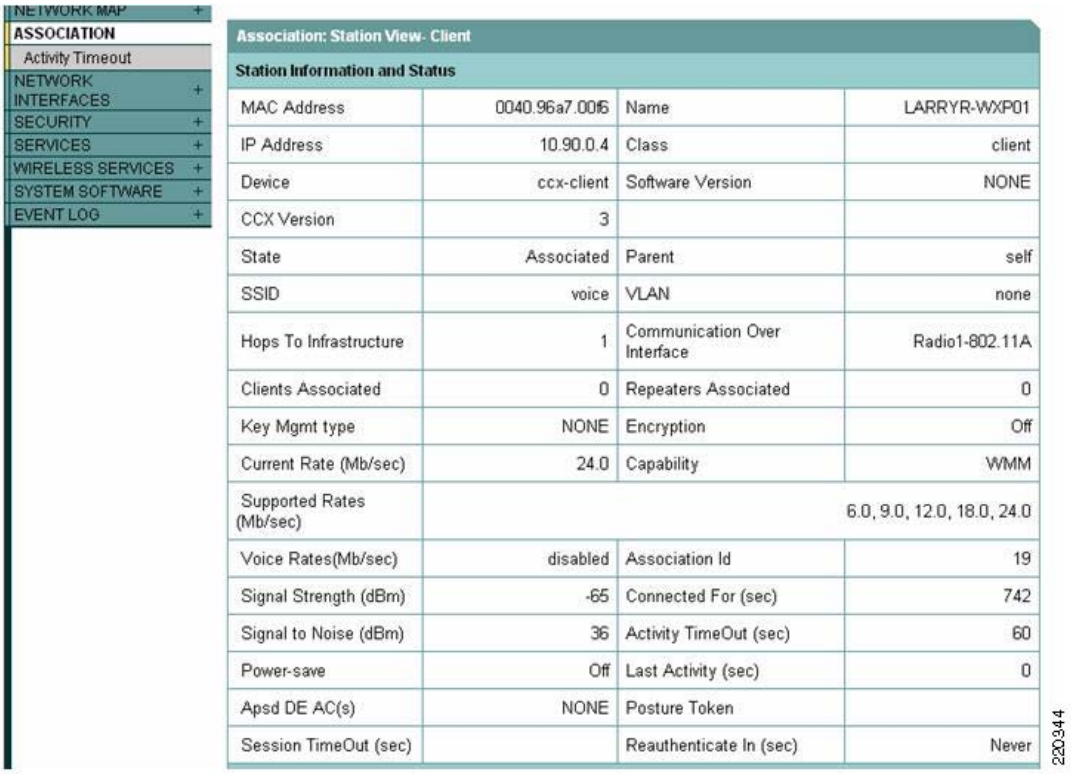

Figure 11-9 WLC Call Metrics

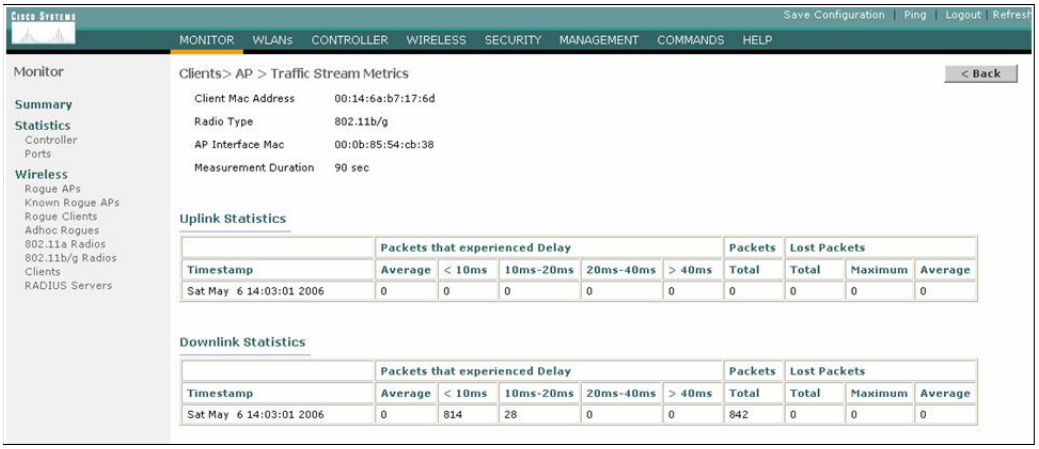

A decoded RTP packet is shown in Figure 11-10. The packet is originated at a 7960 phone. The over-the-air QoS marking is changed from the AVVID marking to a user priority of 6 following the 802.11e specification. Call statistics on the Cisco 7920 and 7921 phones can be viewed on the phone or by browsing into the phone using the IP address of the phone. After that cell edge is determined from the testing of the actual phone, those numbers can then be adapted to more automated tools to complete the coverage design for the site.

Figure 11-10 Sample VoWLAN Capture

When there is multipath interference at the location where dBm measurements are being taken, it is quite likely that the reported dBm values will be different from packet to packet. A packet may be as much as 5dB higher or lower than the previous packet. It may take several minutes to get an average of the signal value at that measuring spot.

AP Call Capacity

A key part of the planning process for a VoWLAN deployment is to plan the number of simultaneous voice streams per AP. When planning the voice stream capacity of the AP, consider the following points:

Note A call between two phones associated to the same AP counts as two active voice streams.

- The utilization of the unlicensed and shared 802.11 channel is the real determinate for the number of simultaneous voice streams an AP carries.
- Because the channel utilization and then the AP performance determine the number of voice streams, same channel and next channel separation are most important. Two APs in the same location and configured to the same channel do not provide double the number of voice streams. In fact, there can be fewer voice streams than one AP would provide.
- Cell capacity or bandwidth determines the number of voice streams that can be simultaneously conducted.
- The handset QoS features supported in the handsets and VoWLAN deployment should be considered.

Various handsets have different WLAN QoS features and capabilities that impact the features that are enabled in the WLAN deployment, and ultimately determine the per-AP call capacity of the AP. Most VoWLAN handsets provide guidance on the number of calls per AP supported by that phone; this should be considered a best case figure where the handset is able to use its optimal QoS features and has full access to the channel capacity.

The actual number of voice streams a channel supports is highly dependent on a number of issues, including environmental factors and client compliance to WMM and CCX specifications. The CCX specifications that are most beneficial to call quality and channel capacity are shown in the table in Figure 11-11. Simulations indicate that a 5 GHz channel can support 14–18 calls. This means a coverage cell can include 20 APs on different channels, with each channel supporting 14 voice streams. The coverage cell can support 280 calls. The number of voice streams supported on a channel with 11b clients is seven; therefore, the coverage cell with three APs on the three non-overlapping channels supports 21 voice streams.

Figure 11-11 CCX VoWLAN Features

How CCX Benefits VoWLAN Call Quality	
Feature	**Benefit**
CCKM Support for EAP-Types	Locally Cached Credentials Means Faster Roams
Unscheduled Automatic Power Save Delivery (U-APSD)	More Channel Capacity and Better Battery Life
TSPEC-Based Call Admission Control (CAC)	Managed Call Capacity for Roaming and Emergency Calls
Voice Metrics	Better and More Informed Troubleshooting
Neighbor List	Reduced Client Channel Scanning
Load Balancing	Calls Balanced Between APs
Dynamic Transmit Power Control (DTPC)	Clients Learn a Power to Transmit At
Assisted Roaming	Faster Layer 2 Roams

220352

Figure 11-11 shows the following:

- Cisco Centralized Key Management (CCKM) provides for faster client roaming for Extensible Authentication Protocol (EAP)-authenticated client, which benefits call quality.
- Call Admission Control (CAC) also benefits call quality and can create bandwidth reservation for E911 and roaming calls.
- Assisted Roaming and Neighbor List benefit call quality and battery life.
- Voice Metrics can benefit management.
- Unscheduled Automatic Power Save Delivery (U-APSD) and Dynamic Transmit Power Control (DTPC) benefit battery life
- Load balancing and DTPC benefit call quality.

Several to the CCX features have more than one benefit.

The amount of buffer memory, CPU speed, and radio quality are key factors of the performance of an AP radio. QoS features prioritize the voice and data traffic in the channel. For a further discussion of QoS, see Chapter 5, "Cisco Unified Wireless QoS."

The 802.11e, WMM, and CCX specifications help balance and prevent the overloading of a cell with voice streams. CAC determines whether there is enough channel capacity to start a call; if not, the phone may scan for another channel. The primary benefit of U-ASPD is the saving of WLAN client power by allowing the transmission of frames from the WLAN client to trigger the forwarding of data frames for a client that has been buffered at the AP for power saving purposes. The Neighbor List option provides the phone a list that includes the channel numbers and channel capacity of neighboring APs. This is done to improve call quality, provide faster roams, and improve battery life.

For a further discussion of U-APSD and CAC, see Chapter 3, "WLAN Radio Frequency Design Considerations."

Cell Edge Design

Guidelines for deploying 802.11b VoWLAN handsets have recommended a design with a minimum power of -67 dBm on the cell boundary (see Figure 11-12). These cells are smaller cells than those used in data WLANs of the past.

The -67 dBm is a general measurement. Achieving a packet error of one percent also requires an SNR value of 25 dB or more. Therefore, when determining the likely channel coverage area for a particular phone type, both signal strength of the phone and the noise must be checked on the AP. See Figure 11-8 and Figure 11-11 for determining these values on the autonomous and LWAPP APs.

The -67 dBm measurement has been used for years for 11b phone clients from many vendors. Tests indicate that this same rule of thumb measurement works well for 11g and 11a phone clients.

Figure 11-12 Cell Edge Measurements

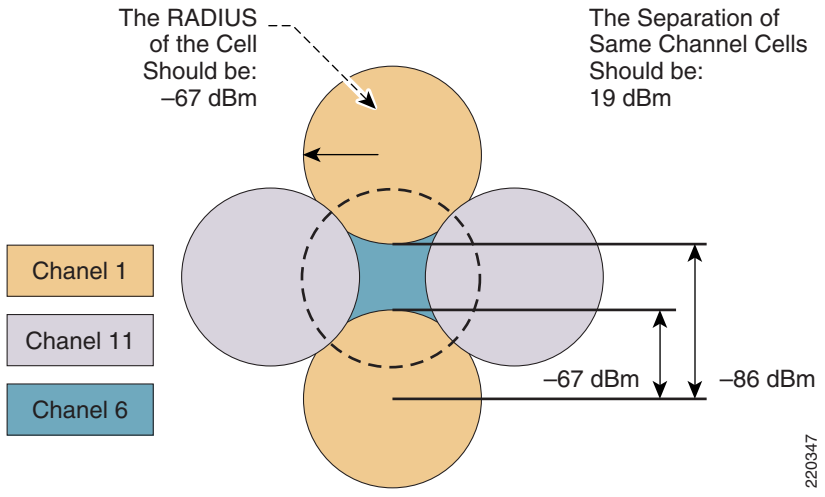

For 5 GHz cells, there is less concern about same channel separation because of the number of channels. There are 20 channels, so a two-channel separation is almost always possible. However, in the 2.4 GHz band, only three channels do not overlay in frequency.

For both 5 GHz and 2.4 GHz, the cell needs to be at the floor location where a packet error rate of 1 percent is maintained at the highest data rate desired for a given channel. In the case of 11b, that data rate is 11 Mbps. Thus, from the center (the AP location) to a point on the floor where the phone signal is seen by the AP, the cell edge is -67 dBm.

802.11g and 802.11a phone clients may be capable of rates up to 54 Mbps. Current chip sets support 54 Mbps, but transmit powers do differ. Cisco highly recommends that all links between phone clients and APs be created with matching transmit powers (see Dynamic Transmit Power Control, page 11-14).

Coverage cells can be created for specific data rates. For a high density deployment or a deployment where a large number of calls are required in a small floor space, 11a is recommended because of the number of channels and the 54 Mbps data rate. The lower data rates in 11a can be disabled, the 24 Mbps data rate can be set to “required”, and the rates of 36 to 54 can be left enabled.

After using the model of cell edge set to -67 dBm, determine where the error rate of 1 percent is, and then examine the SNR value.

Create the phone test to find the –67dB edge by doing the following:

- Set the phone to its desired transmit power.
- Set the AP to a matching transmit power.
- Place the AP and the desired antenna in the location where the phone will be used.
- With an active call, or sending and receiving packets equal in size to the G711 codec, measure the signal level to the –67 dB cell edge.

Carefully examine the data sheets of the particular phone device to determine the transmit powers and data rates supported by the phone device in a particular WiFi band. The data sheets for Cisco Unified Wireless IP Phones can be found on http://www.cisco.com. For phones from other vendors, check the website of the vendor.

The 11a maximum transmit powers vary on different channels and with different AP models. The 11g maximum transmit powers vary by model. Cisco Aironet AP data sheets should be carefully examined to determine which AP model supports which data rates. Figure 11-13 shows an example of the maximum 11a transmit power in dBm by channel.

Figure 11-13 Channel Power Assignment

5GHz	UNII-1				UNII-2				UNII-3			
Channel	36	40	44	48	52	56	60	64	149	153	157	161
Max Tx Power	11	11	11	11	11	17	17	17	17	17	14	11

New UNII-II								
100	104	108	112	116	132	136	140	220353
11	17	17	17	17	17	17	17	

There is a variance of 6 dB from the maximum transmit powers across the 5 GHz band. This means that when using the maximum allowed transmit power throughout a site that allows all channels, there is not equal cell coverage on all channels. It also means that if dynamic channel selection is used, the cell coverage edge may change based on the channel number. However, dynamic channel selection can be tuned (see Chapter 3, “WLAN Radio Frequency Design Considerations.”) The default mode of dynamic channel selection accounts for the difference of maximum transmit powers by channel.

Cell transmit power on all APs should not exceed the maximum or desired transmit power of the phone. If the phones maximum or set transmit power is 13 dBm, Cisco recommends that all APs have a maximum transmit power of 13 dBm. Next, the maximum transmit power on the AP should be set to an equal transmit power or the next higher transmit power. Equal transmit power is recommended to avoid

one-way audio. The AP generally has better receiver sensitivity and diversity support than the phone, so it should be able to receive the slightly lower strength phone signal. See Dynamic Transmit Power Control, page 11-14 for more information on equal transmit powers.

Dual Band Coverage Cells

Chapter 3, "WLAN Radio Frequency Design Considerations," illustrates 2.4 GHz and 5 GHz band channel coverage design. For a dual mode AP to provide equal cell coverage on both the 2.4 GHz channel and the 5 GHz channel, the 2.4 GHz channel must have an equal (or more likely lower) transmit power than the 5 GHz channel.

At most sites, the noise level in the SNR formula will be lower by perhaps 10 dB. The receiver sensitivity of 11g radios is generally 2 dBm better than the same data rate on the 11a radio. As an example, the data sheet for the 7921 has the receive sensitivity of -78 dBm at the data rate of 36 Mbps for 11g, and -76 dBm for 11a. Therefore, given the anticipated better noise floor of 10 dB, the 11a cell can do better by 8 dBm. Other details such as the difference in path loss between 11g and 11a keep this from being a direct ratio. However, if the same coverage cells are desired, reducing the 11g network by one or two power levels from the 11a network should accomplish this goal.

Dynamic Transmit Power Control

Cisco Aironet APs by default have DTPC enabled. DTPC is automatic with the LWAPP controllers and is configurable on the autonomous APs. Clients need CCX version 2 capabilities to use DTPC. DTPC accomplishes the following:

- Sets the phones transmit power to match the transmit power of the AP
- The AP advertises its transmit power for the clients to learn
- Prevents one-way audio; that is, RF traffic is only being heard in one direction

DTPC allows the phone to automatically adjust its transmit power to that of the APs. In the example shown in Figure 11-14, this means that the phone changes its transmit from 5 mW to 100 mW.

Figure 11-14 Client and AP Power Matching

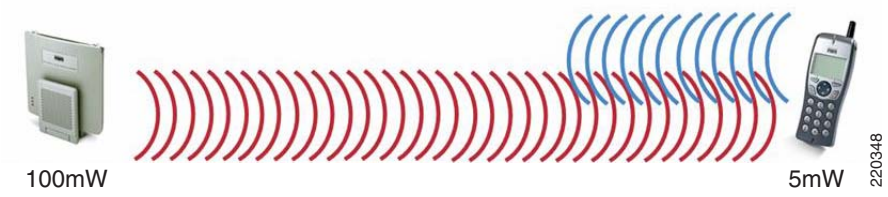

802.11g and 802.11a clients do not have 100 mW transmit powers. Cisco highly recommends that the maximum configured transmit power on the access be no higher than the client phone devices hardware supports. A phone with a slightly lower transmit power than the AP is better than the AP using less power than the phone. Having matching transmit powers lessens the likelihood of one-way audio (the typical user experience of "can you hear me…. I can't hear you").

Interference Sources Local to the User

Interference can be local to the user, but is also likely to affect nearby users. Bluetooth (BT) is a popular RF protocol used in personal area networks that interferes with WiFi 2.4 GHz channels. Figure 11-15 shows that the actual BT signal does span all the 2.4 GHz channels used by 802.11b/g clients. This graphic is from a 802.11g call with a BT headset attached to the phone. Figure 11-15 also shows the jitter caused by the BT headset.

Figure 11-15 Bluetooth (BT) Signal Pattern in the 802.11b/g 2.4 GHz Spectrum of a Typical BT Earpiece

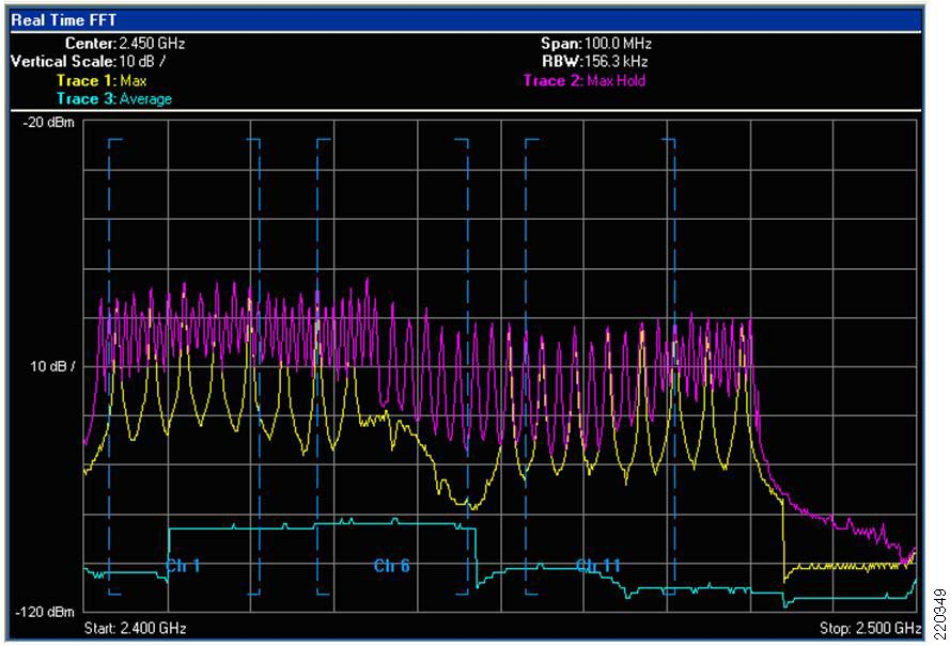

The **PINK** is the Max Hold line, or the line that shows the maximum transmit power that was reached during the test. The **YELLOW** shows the maximum transmit power in the last sample period of ten seconds. The **TURQUOISE** shows the average transmit power over the period of the test. The **vertical dashed** lines separate the three non-overlapping 802.11b/g channels **Ch1, Ch6,** and **Ch11.** The charting is from 2.400 GHz on the left to 2.500 GHz on the right. From the right edge of the Ch11 vertical blue line is the part of the 802.11 spectrum used in Europe and Japan. This capture was done with an AP and clients configured for the North American regulator domain. This graph shows that the BT earpiece was easily transmitting outside of FCC regulations.

Notice that the BT signal is very narrow. BT transmits data on a single MHz of frequency, stops the transmission, moves to another frequency in the 802.11 2.4 GHz band, and then transmits data. This is repeated continually. The 802.11b and 802.11g signals are sent with a combined 22 MHz of frequency. The radio remains on that 22 MHz of frequency. This grouping of 22 MHz is referred to as the channel. The Max Hold line shows how strong the BT is while in search mode. The signal level is above that of a 50 mW (17 dBm) OFDM 802.11g radio. A signal of this strength and duration causes 802.11b/g phones to drop the VoWLAN call. Lesser strength BT signals cause jitter, resulting in a lower MOS value.

Figure 11-16 shows an example of an Ethereal jitter analysis of three simultaneous phone calls, each using a BT earpiece.

Figure 11-16 Jitter Analysis Example

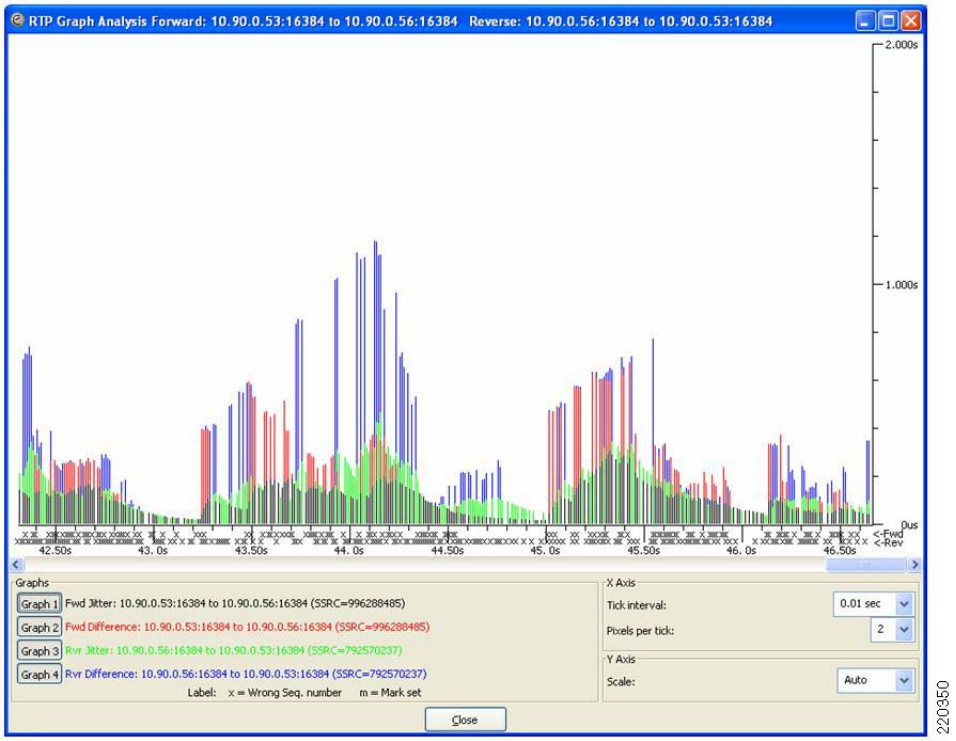

All three calls were on the same AP, and were calls to other phones on this AP.

CHAPTER 12

Cisco Unified Wireless Guest Access Services

The introduction of Wireless LAN (WLAN) technologies in the enterprise has changed the way corporations and small-to-medium businesses function by freeing staff and network resources from the constraints of fixed network connectivity.

WLAN has also changed how individuals access the Internet and their corporate networks from public locations. The advent of Public WLAN (Hotspots) has caused mobile workers to become accustomed to being able to access their corporate network from practically anywhere.

Introduction

The paradigm of public access has extended to the enterprise itself. Long gone is the scenario where it was sufficient for a company to provide its partners, visitors, and guests with a place to sit and possibly an outside line with which to make phone calls. Our highly mobile, information-on-demand culture requires on-demand network connectivity. A half-day spent at a partner or customer venue without access to one's own network resources can impact the productivity of a meeting, service or sales call, and reduce the overall personal productivity of the guest who is away from their office. For this reason, enterprise guest access services are becoming increasingly important and a necessity in the corporate environment.

While there is broad recognition that guest networking is becoming increasingly important, there is also well-founded apprehension over how one safeguards their internal company information and infrastructure assets. Ironically, unbeknownst to many enterprises, their network might already play host to guests who, in an uncontrolled manner, find ways to access the Internet via improperly implemented wired or wireless networks. These guests are not hackers in the true sense, but otherwise well-intentioned individuals trying to get their jobs done. So, on the surface, while it might sound risky to implement a guest access solution, when implemented correctly, an enterprise that implements a guest access solution will most likely improve their overall security posture as a result of the network audits associated with the implementation process.

In addition to overall improved security, implementing a guest access network offers these additional general benefits:

- Authentication and authorization control of guests based on variables including date, duration, and bandwidth
- An audit mechanism to track who is currently using, or has used, the network

Additional benefits of a wireless-based guest access include the following:

- It provides wider coverage by including areas such as lobbies and other common areas that otherwise might not have been wired for network connectivity.

- It removes the need for designated guest access areas or rooms.

Scope

Several architectures can be implemented to offer guest access in the enterprise. It is not the goal of this chapter to cover all possible solutions. Instead, this chapter focuses on the implementation of wireless guest networking using the Cisco centralized controller and lightweight AP (LWAPP) architecture. For more information on deploying wired and wireless Guest Access services in other topology scenarios, see the following URL:
http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html#anchor5

Wireless Guest Access Overview

Ideally, the implementation of a wireless guest network uses as much of an enterprise's existing wireless and wired infrastructure as possible to avoid the cost and complexity of building a physical overlay network. Assuming this is the case, the following additional elements and functions are needed:

- A dedicated guest WLAN/SSID—Implemented throughout the campus wireless network wherever guest access is needed
- Guest traffic segregation—Requires implementing Layer 2 or Layer 3 techniques across the campus network to restrict where guests are allowed to go.
- Access control—Involves using imbedded access control functionality within the campus network or implementing an external platform to control guest access to the Internet from the enterprise network.
- Guest user credential management—A process by which a sponsor can create temporary credentials in behalf of a guest. This function might be resident within an access control platform or it might be a component of AAA or some other management system.

Wireless Guest Access using a Centralized Controller Architecture

The Cisco Centralized WLAN solution offers a flexible, easy-to-implement method for deploying wireless guest access by using Ethernet in IP (RFC3378) within the centralized architecture. Ethernet in IP is used to create a tunnel across a Layer 3 topology between two WLAN controller endpoints. The benefit of this approach is that there are no additional protocols or segmentation techniques that must be implemented to isolate guest traffic from the enterprise. See Figure 12-1 for an example of guest access topology using a centralized WLAN architecture.

Figure 12-1 Centralized Controller Guest Access

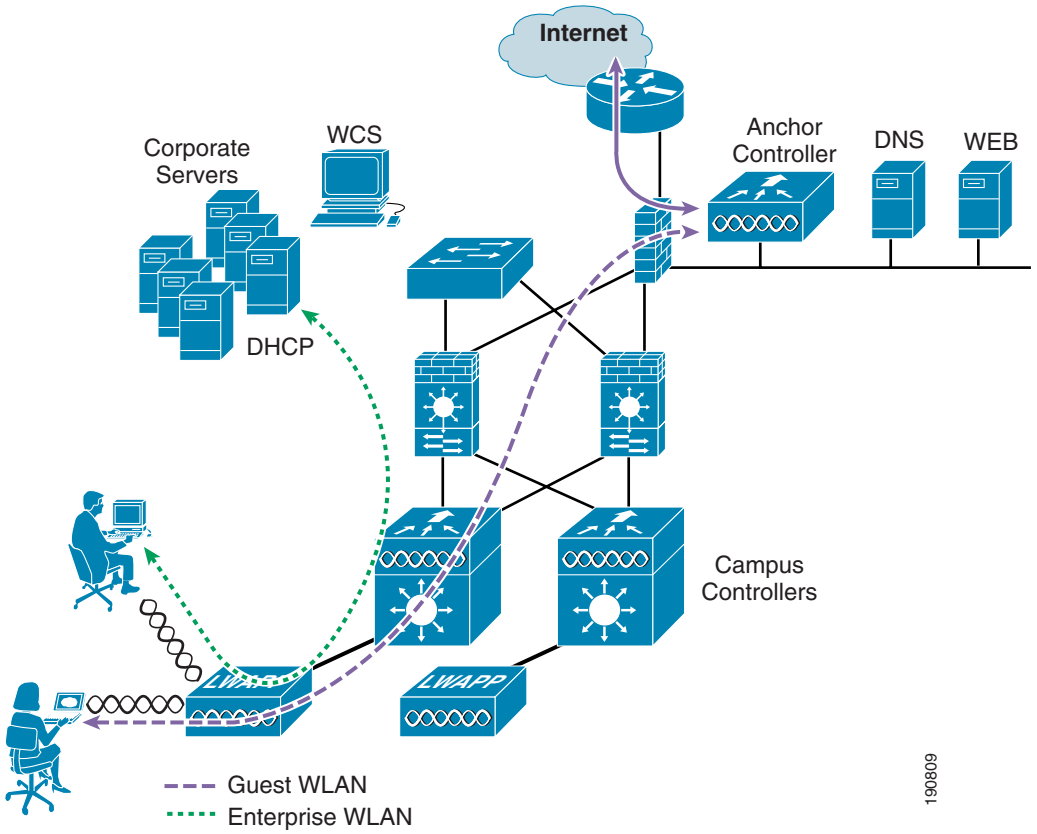

As shown in Figure 12-1, a controller is located in the enterprise's DMZ where it performs an anchor function. This anchor controller is responsible for terminating EoIP tunnels that originate from other campus WLAN controllers throughout the network. Remote controllers are responsible for termination, management, and standard operation of the various WLANs provisioned throughout the enterprise, including one or more guest WLANs. Guest WLANs, instead of being bridged locally to a corresponding VLAN, are bridged via an EoIP tunnel to the anchor controller. Specifically, guest WLAN data frames are bridged via LWAPP to the remote controller and via EoIP for the campus WLC to a guest VLAN defined at the anchor WLC. In this way, guest user traffic is forwarded to the Internet transparently, with no visibility by, or interaction with, other traffic in the enterprise.

Non-Controller Based Wireless Guest Access

In place of using a centralized WLAN architecture, the biggest challenge in implementing guest access services is the segmentation (isolation) of guest traffic from the rest of the enterprise. This is especially true for wired networks or wireless networks that use autonomous (fat) APs. Some method of traffic segmentation must be implemented beginning with a separate WLAN or VLAN, coupled with a policy that is applied at the first Layer 3 hop in the network. Possible options include the following:

- Distributed ACLs—This method involves implementing ACLs throughout the enterprise to restrict guests from accessing network resources within the host enterprise network (see Figure 12-2).

Figure 12-2 Segmentation using Distributed ACLs

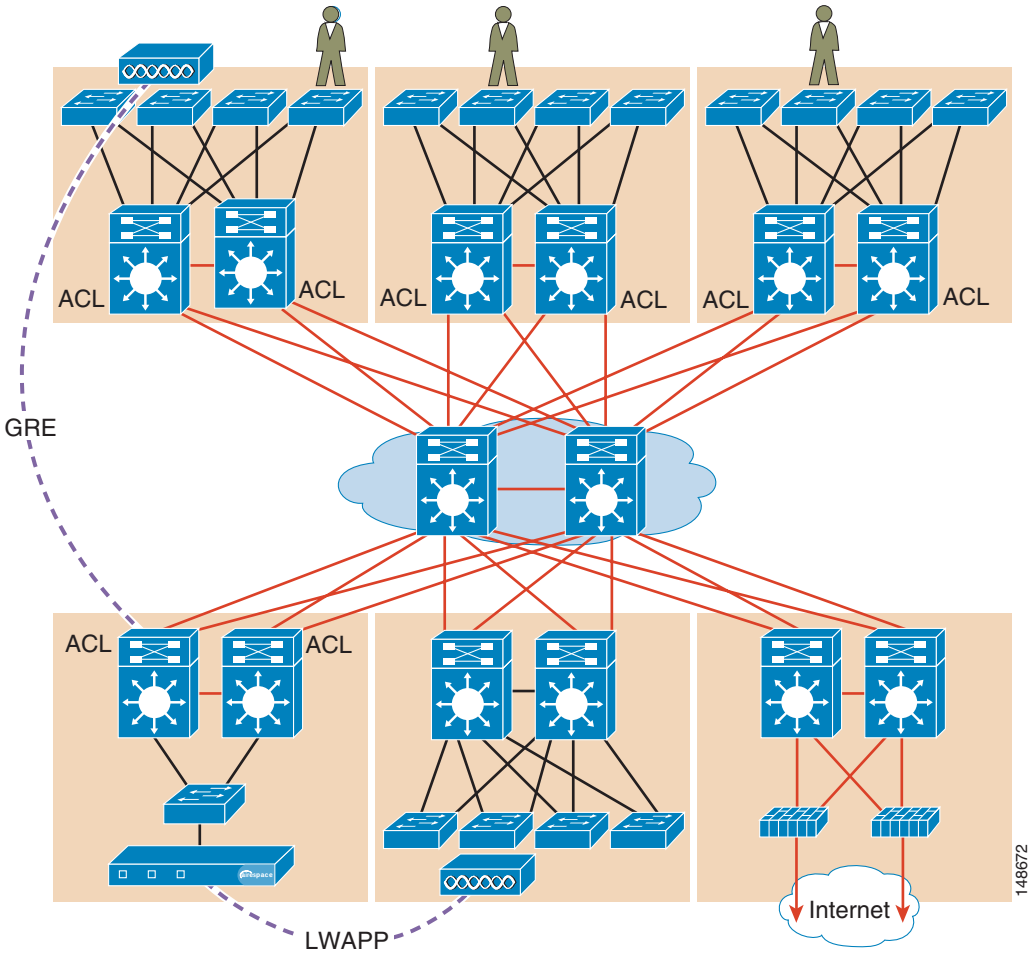

- Policy based routing (PBR) and GRE tunnels—GRE tunnels are used to create a logical overlay network through which guest traffic is directed to the Internet edge or DMZ. Policy-based routing is used to classify and enforce guest traffic into the tunnels. Companies that have deployed, or are looking to deploy, a PBR/GRE should seriously consider, if at all possible, VRF-lite with GRE (see Figure 12-3).

Figure 12-3 Segmentation using PBR into GRE Tunnels

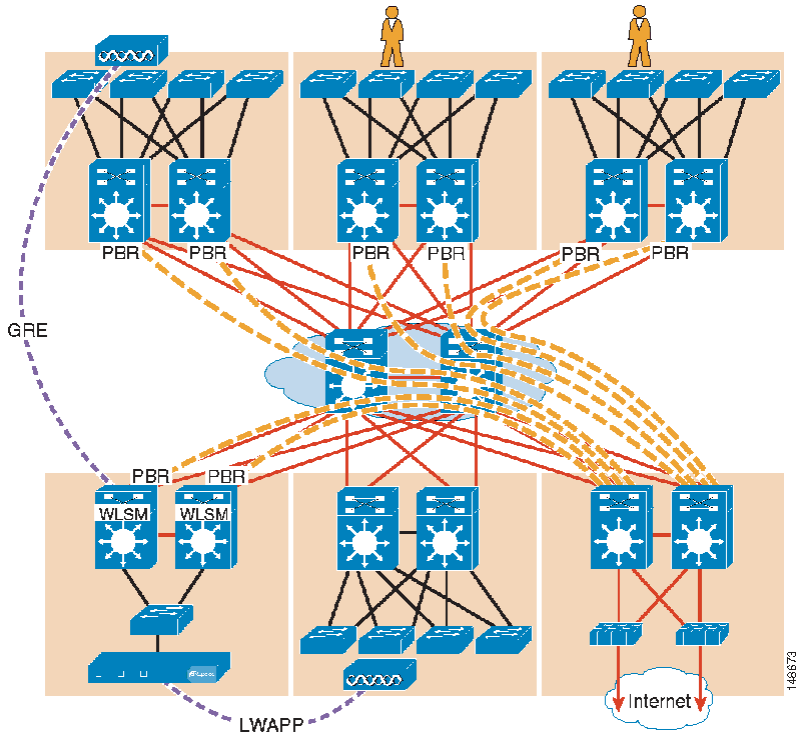

- VRF-Lite and GRE or mGRE tunnels—This technique is similar to the PBR/GRE method. An overlay network is created through the implementation of GRE tunnels. However, instead of policy routing traffic into the tunnels, the guest access VLAN and interfaces, along with the tunnel interfaces, are assigned to a virtual routing and forwarding (VRF) instance. This ensures guest traffic can be forwarded only through the tunnels (see Figure 12-4 and Figure 12-5).

Figure 12-4 Segmentation using VRF-lite and GRE

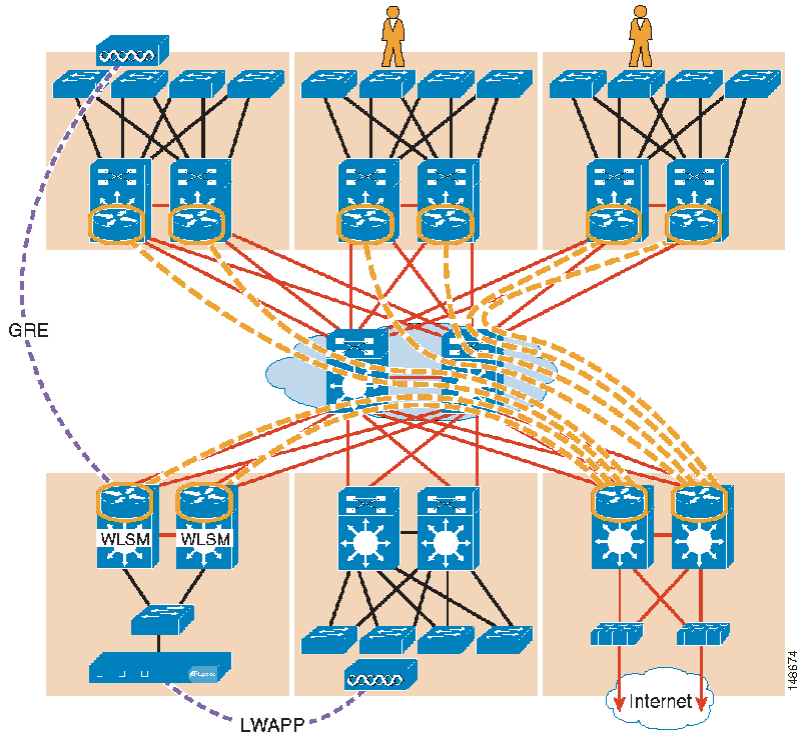

Figure 12-5 Segmentation with VRF-Lite and mGRE

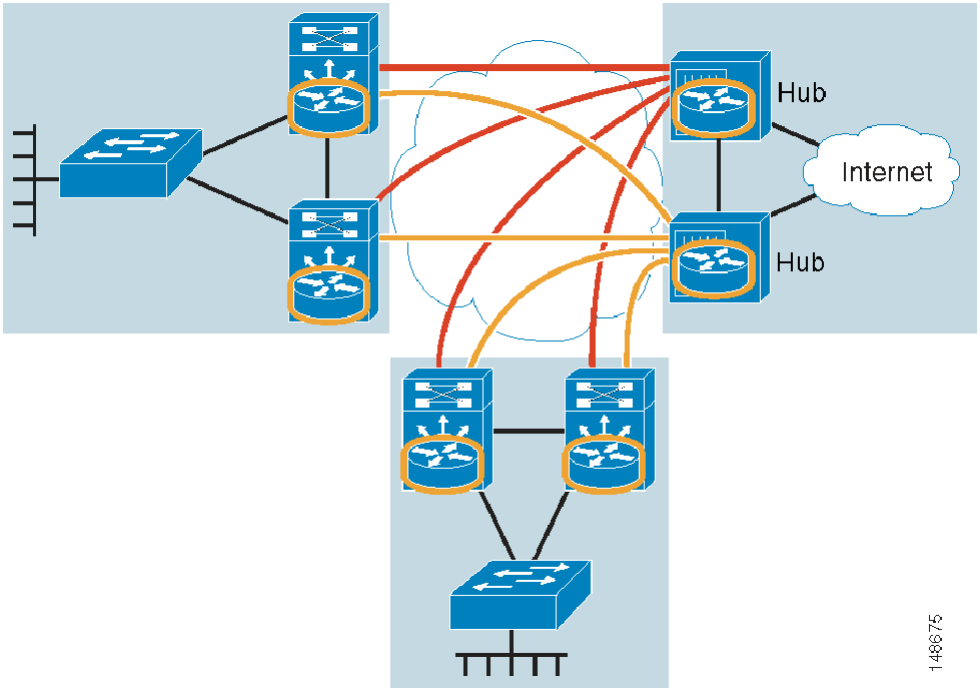

Each of the methods described here has benefits and drawbacks. Additional challenges for the enterprise include the following:

- Determining which method to use
- Modifying existing topologies to accommodate a given segmentation method or introducing new versions of code into the network
- The added complexity of managing a logical overlay network

In addition, it simply might not be practical or possible for an enterprise to implement one of the segmentation methods described in this chapter. In that case, and if wireless guest access is all that is needed, deploying guest services using a centralized WLAN controller architecture is a good alternative.

Otherwise, if an enterprise requires both wireless and wired guest access, see the guest access documentation at the following URL:

http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidances_list.html#anchor5

The remainder of this chapter focuses on the implementation of wireless guest networking using the Cisco Centralized Controller solution.

Wireless Controller Guest Access

The Cisco Wireless Controller Guest Access solution is self-contained and does not require any external platforms to perform access control, web portal, or AAA services. All these functions are configured and run within the anchor controller. However, the option exists to implement one or all of these functions externally and that is covered later in the chapter.

Supported Platforms

The anchor function, which includes tunnel termination, web authentication, and access control is supported on the following WLC platforms (using version 4.0 and later software images):

- Cisco 4400 Series
- Cisco 6500 Series (WISM)
- Cisco 3750 with integrated WLC

The following WLC platforms cannot be used for anchor functions, but can be used for normal controller operations and guest tunnel origination to an anchor controller:

- Cisco WLAN Controller Module for Integrated Service Routers (ISR)
- Cisco 2000 Series

WLAN Anchors and Ethernet in IP to Support Guest Access

A key feature of the Cisco centralized controller architecture is the ability to statically map one or more provisioned WLANs to a specific controller (anchor) within the network, using an EoIP tunnel. By using this technique, a guest WLAN and all associated guest traffic can be transported transparently across the enterprise network to an anchor controller that resides in the Internet DMZ (see Figure 12-6).

Figure 12-6 Static EoIP Tunnels

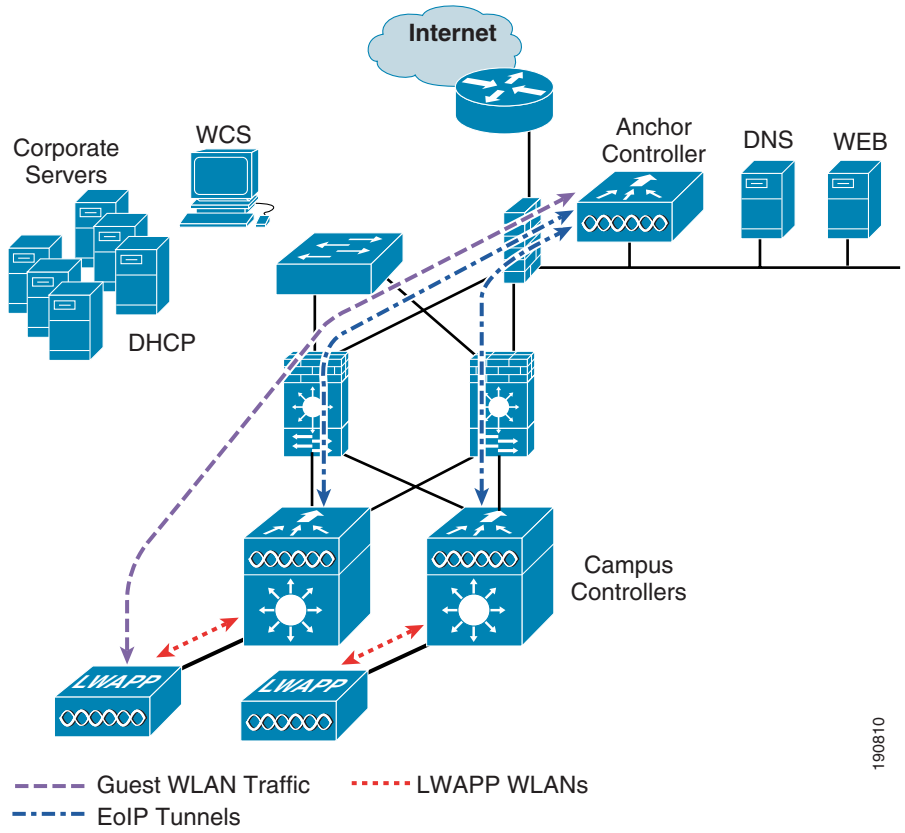

Figure 12-7 shows a sniffer trace of an Ethernet in IP tunnel (highlighted) between a branch wireless controller with a guest WLAN provisioned and an anchor controller that is performing local web authentication. The first IP detail shown represents the Ethernet in IP tunnel between the branch and anchor controllers. The second IP detail is that of guest traffic (in this case, a DNS query).

Figure 12-7 Sample Ethernet in IP Sniffer Trace

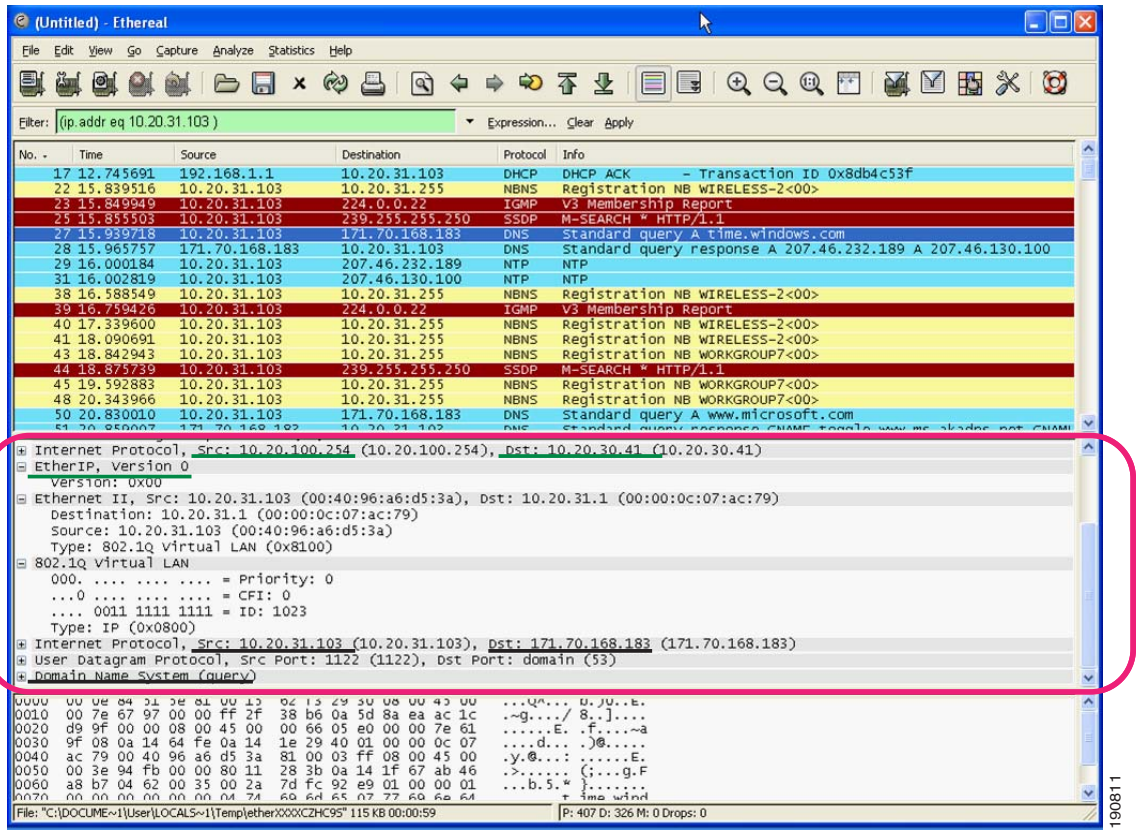

Anchor Controller Deployment Guidelines

This section provides guidelines for deploying an anchor controller to support wireless guest access.

Anchor Controller Positioning

Because the anchor controller is responsible for termination of guest WLAN traffic and subsequent access to the Internet, it should be positioned in the enterprise's Internet DMZ. By doing so, rules can be established within the firewall to precisely manage communications between authorized controllers throughout the enterprise and the anchor controller. Such rules might including filtering on source or destination controller addresses, UDP port 16666 for inter-WLC communication, and IP protocol ID 97 Ethernet in IP for client traffic. Other rules that might be needed include the following:

- TCP 161 and 162 for SNMP
- UDP 69 for TFTP
- TCP 80 or 443 for HTTP, or HTTPS for GUI access
- TCP 23 or 22 for Telnet, or SSH for CLI access

Depending on the topology, the firewall can be used to protect the anchor controller from outside threats.

For the best possible performance and because of its suggested positioning in the network, it is strongly recommended that the guest anchor controller be dedicated to supporting guest access functions only. In other words, the anchor controller should not be used to support guest access in addition to controlling and managing LWAPP APs in the enterprise.

DHCP Services

As previously described, guest traffic is transported at Layer 2 via EoIP. Therefore, the first point at which DHCP services can be implemented is either locally on the anchor controller or the controller can relay client DHCP requests to an external server. See Guest Access Configuration, page 12-16 for configuration examples.

Routing

Guest traffic egress occurs at the anchor controller. Guest WLANs are mapped to a specific physical interface or VLAN on the anchor. Depending on the topology, this interface or VLAN might connect to an interface on a firewall, or directly to an Internet border router. Therefore, a client's default gateway IP is either that of the firewall or the address of a VLAN or interface on the first hop router. For ingress routing, it is assumed the guest VLAN is directly connected to a DMZ interface on a firewall or to an interface on a border router. In either case, the guest (VLAN) subnet is known as a directly connected network and advertised accordingly.

Anchor Controller Sizing and Scaling

The most cost-effective platform to support guest networking in most enterprise deployments is the Cisco 4400 Series controller. Assuming the controller is being deployed to support guest access and tunnel termination functions only, the 4402 with support for 12 APs is sufficient because it is assumed the controller is not going to be used to manage LWAPP APs in the network.

A single 4400 Series controller can support EoIP tunnels from up to 40 other controllers within the enterprise. Additionally, the 4400 supports up to 2500 simultaneous users and has a forwarding capacity of 2 Gbps.

Anchor Controller Redundancy

Anchor controller redundancy is expected to be supported in future releases.

Web Portal Authentication

The Cisco Centralized Guest Access solution can make use of a built-in web portal that is used to solicit guest credentials for authentication and offers simple branding capabilities, along with the ability to display disclaimer or acceptable use policy information (see Figure 12-8).

Figure 12-8 Controller Web Authentication Page

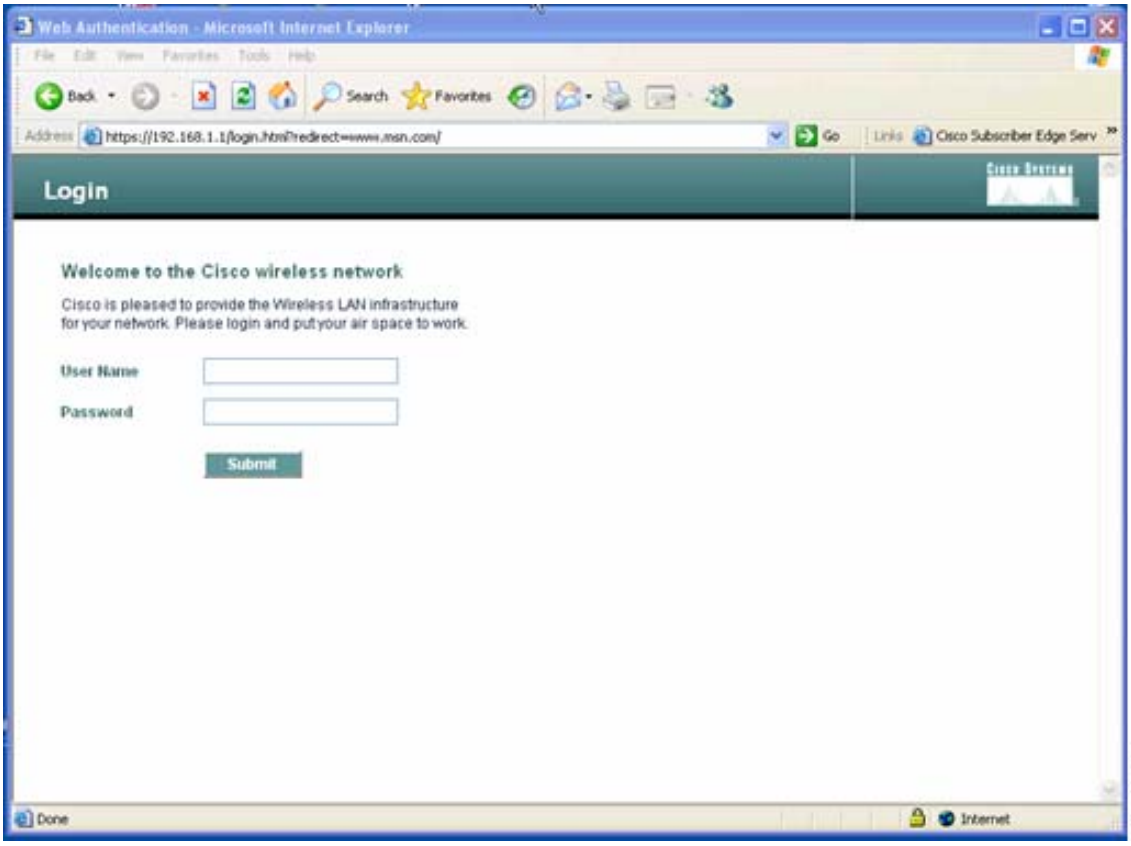

The web portal page is available on all Cisco WLAN controller platforms and is invoked by default when a WLAN is configured for Layer 3 web policy-based authentication.

If a more complex page is required, administrators have the option of importing and locally storing a customized page. Additionally, if an enterprise wants to use an external web server, the controller can be configured to redirect to it in place of using the internal server. See Guest Access Configuration, page 12-16 for web page configuration guidelines.

User Redirection

As is typical for most web-based authentication access, for guest clients to be redirected to the controller web authentication page, they must launch a web browser session and attempt to open a destination URL. For redirection to work correctly, the following conditions must be met:

- DNS resolution—The guest access topology must ensure that valid DNS servers are assigned via DHCP and those DNS servers are reachable to users before authentication. When a client associates to a web policy WLAN for authentication, all traffic is blocked except DHCP and DNS. Therefore, DNS servers must be reachable from the anchor controller. Depending on the topology, this might require opening up conduits through a firewall to permit DNS or modifying ACLs on an Internet border router.

Note Clients with static DNS configurations might not work depending on whether the specified DNS servers are reachable from the guest network.

- Resolvable Home Page URL—A guest user home page URL must be globally resolvable by DNS. If a user home page is, for example, an internal company home page that cannot be resolved outside of their company intranet, that user is not redirected. In this case, the user must open a URL to a public site such as www.yahoo.com or www.google.com.
- HTTP Port 80—If a user home page connects to a web server on a port other than HTTP port 80, they are not redirected. Again, the user needs to open a URL that uses port 80 to be redirected to the controllers web authentication page.

Note In addition to port 80, there is an option to configure one additional port number that the controller can monitor for redirection.The setting is available only through the CLI of the controller: <controller_name> config> **network web-auth-port** <port>.

Guest Credentials Management

Guest credentials can be created and managed centrally using WCS beginning with release 4.0 and later. A network administrator can establish a limited privilege administrative account within WCS that permits lobby ambassador access for the purpose of creating guest credentials. The only function a lobby ambassador is permitted to do is create and assign guest credentials to a controller serving as an anchor for wireless guest access. See Guest Access Configuration, page 12-16 for configuration guidelines.

As with many configuration tasks within WCS, guest credentials are created using templates. For more information regarding administration and use of WCS, see Chapter 8, “Cisco Unified Wireless Control System.” A guest user template includes the following information:

- User name
- Auto generate password (check box) or Administrator assigned password
- Confirm password
- SSID (select box)—Only WLANs configured for Layer 3 web policy authentication are displayed
- Description
- Credentials lifetime—days:hours:minutes

After a lobby ambassador has created a guest template, it can be applied to one or more controllers. Only controllers with web policy-configured WLANs are listed as a candidate controller to which the template can be applied.

Guest credentials, when applied, are stored locally on the anchor controller and remain there until expiration of the Lifetime variable defined in the guest template. If a wireless guest is associated and active when their credentials expire, the WLAN controller stops passing traffic from that user, causing them to be disconnected. Unless the guest credentials are re-applied (to the controller), the user is no longer able to authenticate and access the network.

Note The Lifetime variable associated with guest credentials is independent of the WLAN session time-out variable. If a user remains connected beyond the guest WLAN session time-out interval, they are de-authenticated. The user is then redirected to the web portal and, assuming their credentials have not expired, is required to log back in. To avoid annoying redirects for authentication, the guest WLAN session time-out variable should be set accordingly.

Local Controller Lobby Admin Access

In the event a centralized WCS management system is not deployed, a network administrator can establish a local user management account on the anchor controller, which has lobby admin privileges only. A person who logs in to the controller using the lobby admin account has access only to guest user management functions. Options available for local guest management are the same as for guest template creation in WCS:

- User name
- Auto generate password (check box) or administrator-assigned password
- Confirm the password
- SSID (check box)

  Only WLAs configured for Layer 3 web policy authentication are displayed.
- Description
- Credentials lifetime—days:hours:minutes

Any credentials that may have been applied to the controller by WCS are shown when an admin logs into the controller using the lobby admin account. The lobby admin account has rights to modify or delete any guest credentials that were created by WCS.

Guest User Authentication

As was previously covered in Guest Credentials Management, page 12-12, when an administrator uses WCS or a local admin account on a controller to create guest user credentials, those credentials are stored locally on the controller, which in the case of a centralized guest access topology, would be the anchor controller.

When a wireless guest logs in through the web portal, the controller handles the authentication in the following order:

1. The controller checks its local database for username and password and, if present, grants access.

If no user credentials are found, then:

2. The controller checks to see if an external RADIUS server has been configured for the guest WLAN (under WLAN configuration settings). See External Radius Authentication, page 12-38 for a configuration example. If so, then the controller creates a RADIUS access-request packet with the user name and password and forwards it to the selected RADIUS server for authentication.

If no specific RADIUS servers have been configured for the WLAN:

3. The controller checks its global RADIUS server configuration settings. Any external RADIUS servers configured with the option to authenticate “network” users are queried with the guest user credentials. See External Radius Authentication, page 12-38 for a configuration example. Otherwise, if no servers have network user selected, and the user has not authenticated as a result of 1 or 2 above, the authentication fails.

Note A RADIUS server can still be used to support network user authentication even if the network user check box is cleared under the global settings. However, to do so, that server must be explicitly selected under a given WLANs RADIUS server settings. See External Radius Authentication, page 12-38 for a configuration example.

External Authentication

The guest account creation (Lobby Ambassador) capabilities available in WCS and locally on the controller can be used only to create and store guest user information locally on the controller. An enterprise may already have developed similar functionality in conjunction with a wired guest access or NAC-type deployment. In this case, the anchor controller/guest WLAN can be configured to forward web portal authentication to an external RADIUS server, as described in Guest User Authentication, page 12-13.

The default protocol used by the controller to authenticate web users is Password Authentication Protocol (PAP). In the event you are authenticating web users to an external AAA server, be sure to verify the protocols supported by that server. The anchor controller can also be configured to use CHAP or MD5-CHAP for web authentication. The method is configured under the general configuration settings of the controller by using the web administrative interface of the controller, or through WCS.

External Authentication using Cisco Secure ACS and Microsoft User Databases

If a guest access deployment is planning to use a Microsoft user database in conjunction with Cisco ACS to authenticate guest users, see the following additional Cisco ACS configuration caveats:
http://www.cisco.com/univercd/cc/td/doc/product/access/acs_soft/csacs4nt/acs40/install/postin.htm

See specifically the following URL:
http://www.cisco.com/univercd/cc/td/doc/product/access/acs_soft/csacs4nt/acs40/install/postin.htm#wp1041223

Guest Pass-through

Another variation of wireless guest access is to bypass user authentication altogether and allow open access. However, an enterprise might still want to present an acceptable use policy or disclaimer page to users before granting access. If this is the case, then a guest WLAN can be configured for web policy pass through. In this scenario, a guest user is redirected to a portal page containing disclaimer information. Pass through mode also has an option for a user to enter an e-mail address before connecting (see Figure 12-9 and Figure 12-10 for sample pages). See Guest Access Configuration, page 12-16 for configuration examples.

Figure 12-9 Pass-through Welcome AUP Page

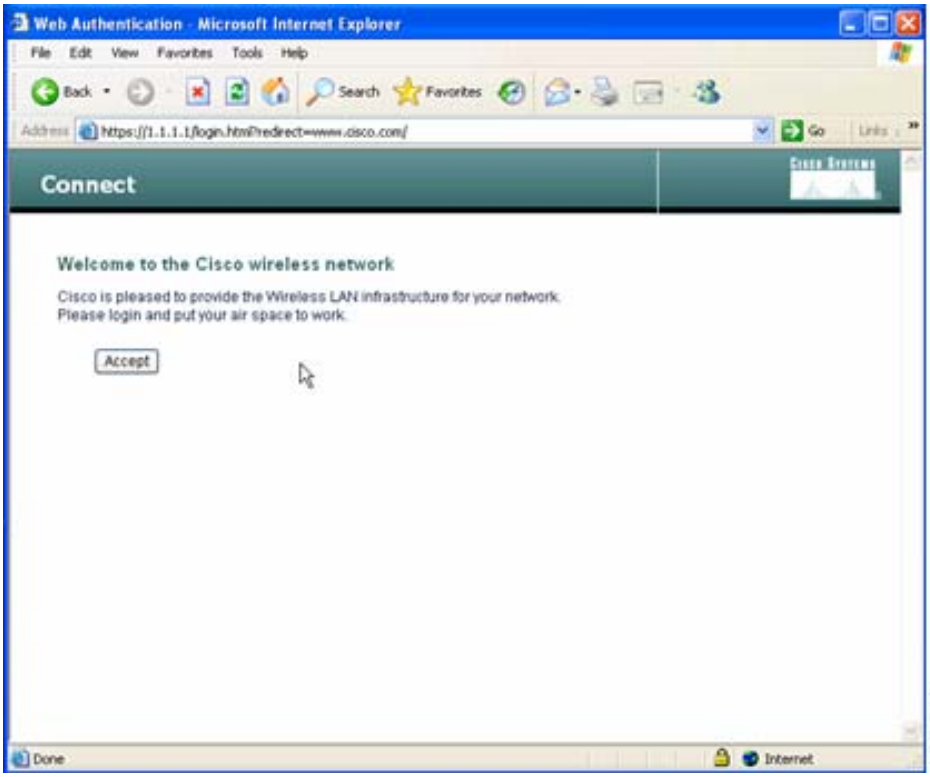

Figure 12-10 Pass-through Page with Email

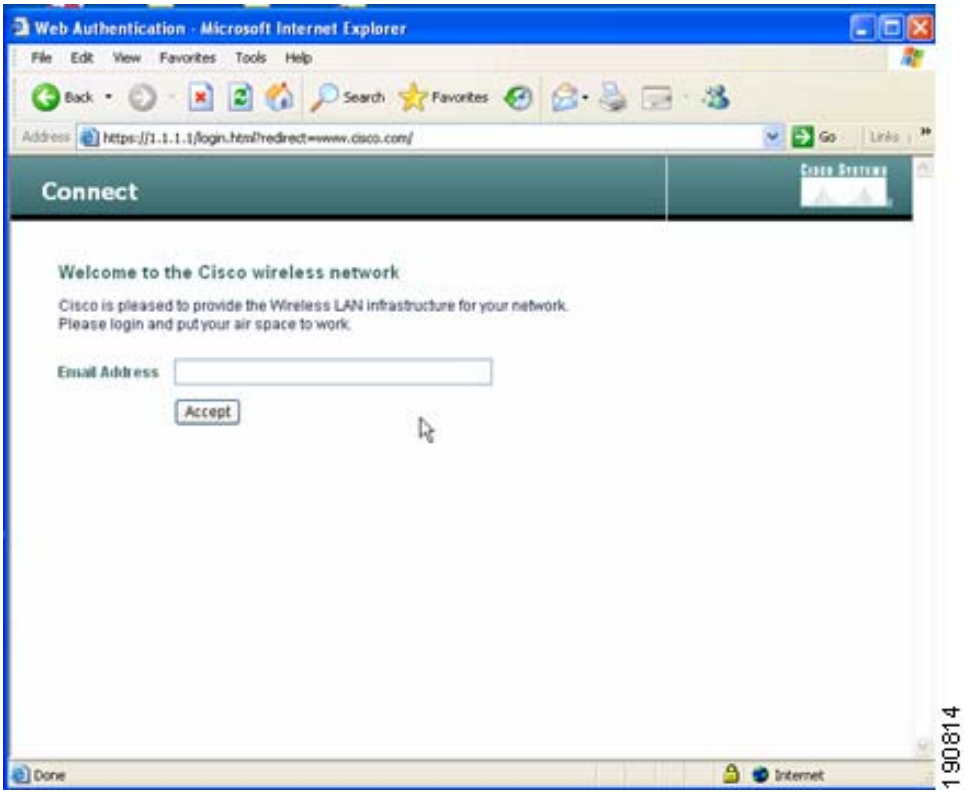

Guest Access Configuration

This section describes how to enable a wireless guest access service using the Cisco Centralized WLAN Controller solution. The configuration tasks require the use of a web browser, Windows IE6 (only). A web session is established with the controller by opening an HTTPS session to the controller management IP address: **https://***management_IP* or optionally to a controller service port IP address.

The following procedures assume there is already a deployed infrastructure of controllers and lightweight APs with the possible exception of the anchor controller platform. See Anchor Controller Deployment Guidelines, page 12-9 for more information.

Note Cisco recommends that the configuration steps outlined in this section are followed in the order in which they are presented.

The following references are used throughout the configuration sections:

- Remote Controller—Refers to the one or more controllers deployed in an enterprise campus or branch location that are used for managing and controlling a group of lightweight APs. Remote controllers map a guest WLAN into a static EoIP tunnel.
- Anchor Controller—Refers to a controller deployed in the enterprise DMZ that is used to perform EoIP tunnel termination, web redirection, and user authentication.

Note Only the relevant portion of a configuration screen capture is shown in this section.

Anchor Controller Interface Configuration

As described in Anchor Controller Positioning, page 12-9, Cisco strongly recommends that the anchor controller be dedicated to guest access functions and not be used to control and manage lightweight APs in the enterprise.

This section does not address all aspects of interface configuration on the anchor controller. It is assumed the reader is familiar with the controller initialization and configuration process required upon initial boot-up using the serial console interface. If not, see the following URL:

http://www.cisco.com/en/US/customer/products/ps6366/products_quick_start_book09186a008056aaa0.html

This section offers specific information and caveats as it pertains to configuring interfaces on a controller being deployed as an anchor in a guest access topology.

As part of the initial configuration (through serial console interface), you are required to define three static interfaces:

- Controller management—This interface/IP is used for communications with other controllers in the network. It is also the interface used to terminate EoIP tunnels that originate from the remote controllers.
- AP Manager interface—Even though the controller is not used to manage APs, you are still required to configure this interface. The easiest thing to do is to configure the interface to be on the same VLAN and subnet as the primary management interface.

Note This would not otherwise be recommended in a traditional deployment where the controller is managing lightweight APs.

- Virtual Interface—The controller quick start installation documentation recommends defining the virtual IP with an address, such as 1.1.1.1. This address needs to be the same for all controllers that are members of the same mobility group name. This virtual interface is also used as the source IP address when the controller performs web authentication.

Guest VLAN Interface Configuration

The interfaces previously described are for operations and administrative functions associated with the controller. To implement a guest access service, another interface must be defined. This is the interface through which guest traffic is forwarded for routing to the Internet. As previously described in Anchor Controller Positioning, page 12-9, the guest interface will likely connect to a port on a firewall or be switched to an interface on an Internet border router.

Defining a New Interface

Perform the following to define and configure an interface to support guest traffic:

Step 1 Click the **Controller** tab.

Step 2 In the left pane, click **Interfaces**. (See Figure 12-11.)

Figure 12-11 Interfaces

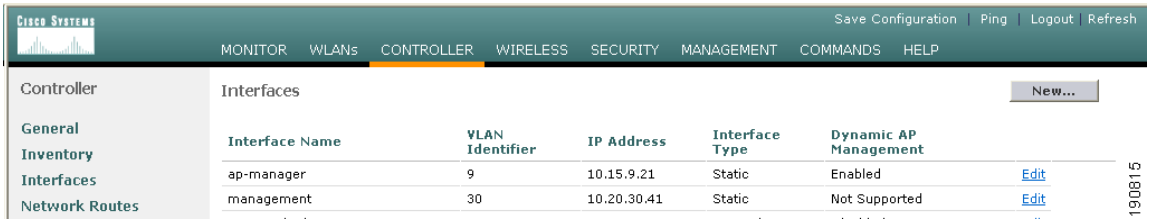

Defining an Interface Name and VLAN ID

Step 3 Enter an interface name and VLAN ID. (See Figure 12-12.)

Figure 12-12 Interface Name and VLAN ID

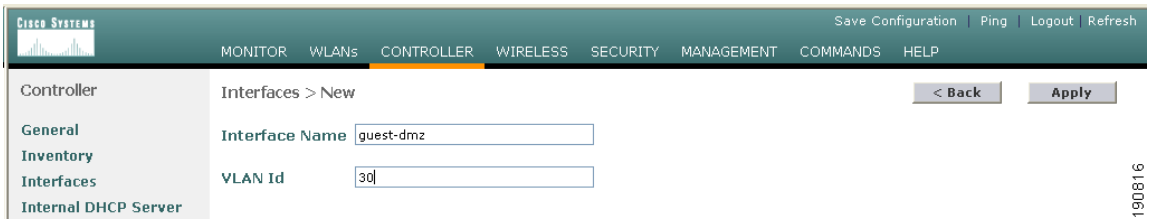

Defining Interface Properties

Step 4 Define the following properties:

- Interface IP
- Mask
- Gateway (for the firewall or next hop router connected to the anchor controller)
- DHCP Server IP (If using an external DHCP server, use the IP address of that server in the Primary DHCP Server field.)

  See Figure 12-13.

Figure 12-13 Defining Properties

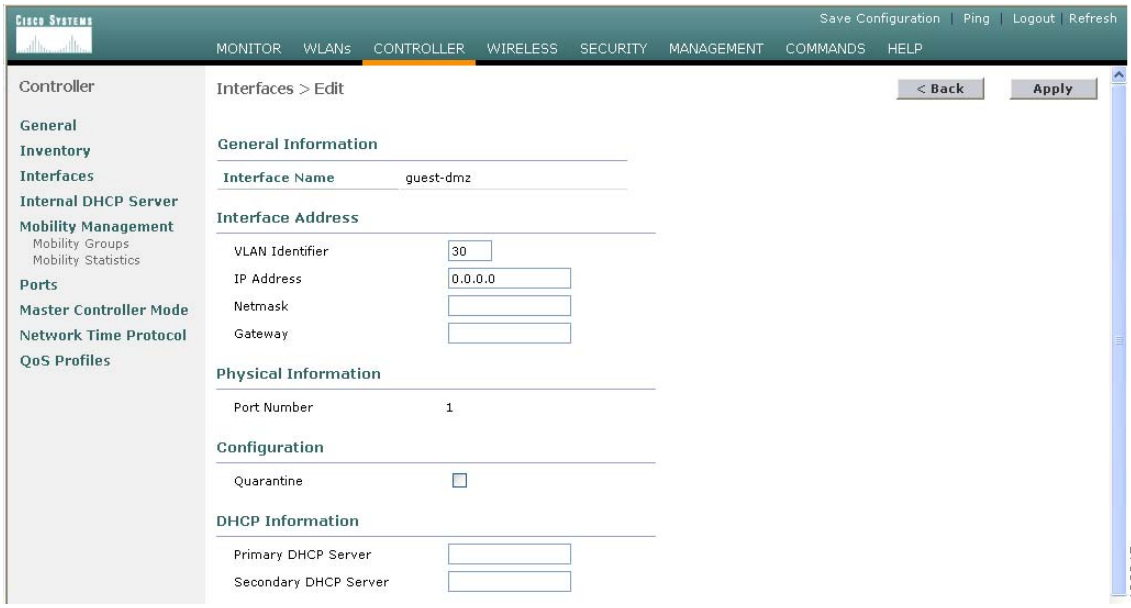

Note If DHCP services are to be managed locally on the anchor controller, populate the primary DHCP server field with the controller management IP address. See Anchor Controller Interface Configuration, page 12-17.

Anchor Controller DHCP Configuration (Optional)

If the anchor controller is going to manage DHCP services for the guest access WLAN, proceed with the following steps.

Adding a New DHCP Scope to the Anchor Controller

Step 1 Click the **Controller** tab.

Step 2 In the left pane, click **Internal DHCP Server**.

Step 3 Click **New**. (See Figure 12-14.)

Figure 12-14 Adding a New DHCP Scope

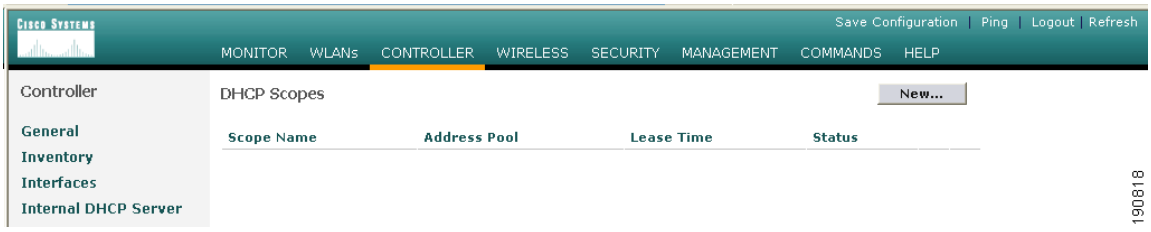

Defining a Scope Name

Step 4 Define a name for the scope and click **Apply**. (See Figure 12-15.)

Figure 12-15 Defining a Scope Name

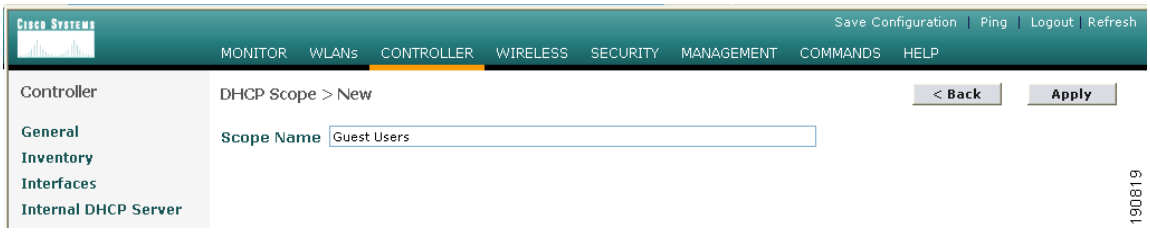

Step 5 Click **Edit**. (See Figure 12-16.)

Figure 12-16 Editing Scope Name

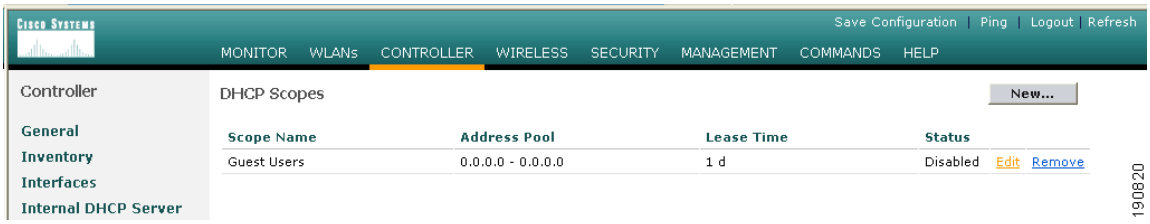

Defining Scope Properties

Step 6 Define the following minimum information:

- Pool start and stop
- Network
- Mask
- Default routers
- DNS servers

Step 7 For Status, choose **Enabled** and click **Apply**. (See Figure 12-17.)

Figure 12-17 Enabling Scope Properties

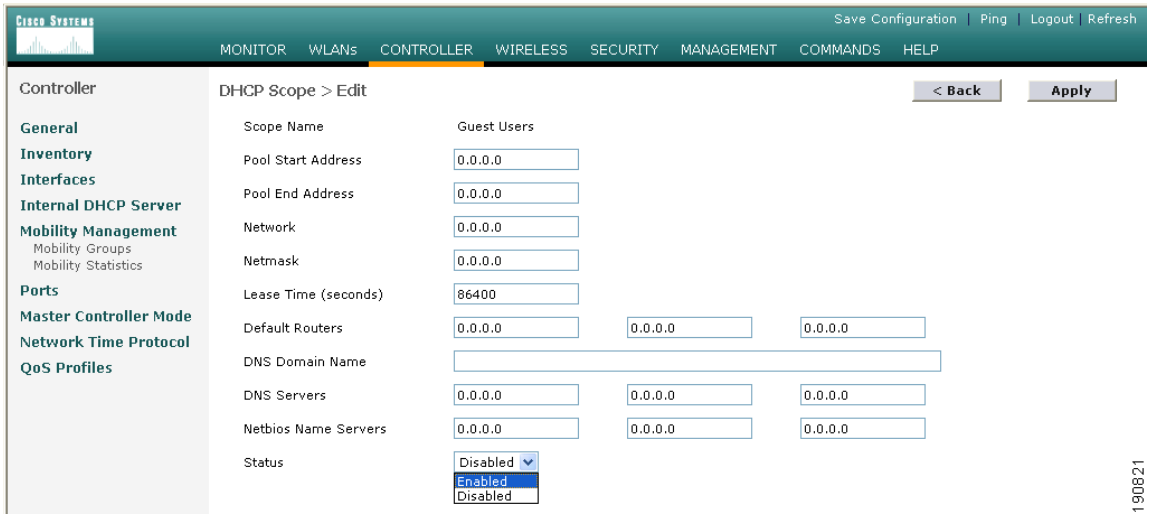

Mobility Group Configuration

The following mobility group parameters should already be defined on the remote controllers as part of a standard centralized WLAN deployment. The anchor controller can also be configured with a mobility domain name, but it is not required to support termination of guest WLAN EoIP tunnels.

Defining a Default Mobility Domain Name for the Anchor Controller (Optional)

Assign a default mobility domain name for the anchor controller.

Step 1 Click the **Controller** tab.

Step 2 Enter a name in the **Default Mobility Domain Name** field.

Step 3 Click **Apply**. (See Figure 12-18.)

Figure 12-18 Defining a Default Mobility Domain Name

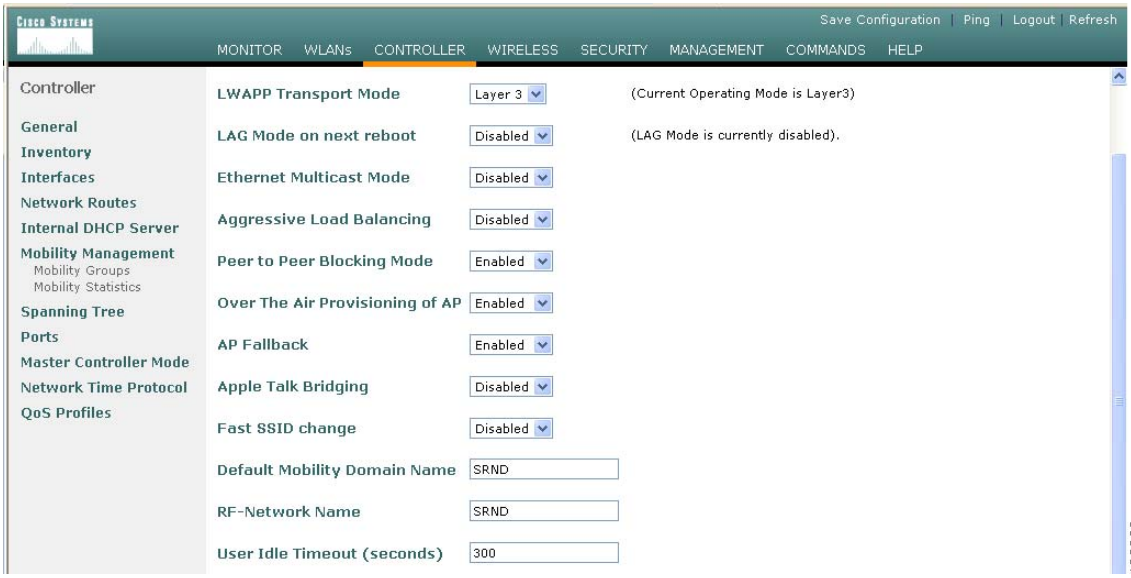

Defining Mobility Group Members for the Anchor Controller

Each campus (foreign) controller in the enterprise that will support the guest WLAN must be defined as a mobility group member in the anchor controller.

Step 1 Click the **Controller** tab.

Step 2 In the left pane, click **Mobility Groups**. (See Figure 12-19.)

Figure 12-19 Defining Mobility Group Members

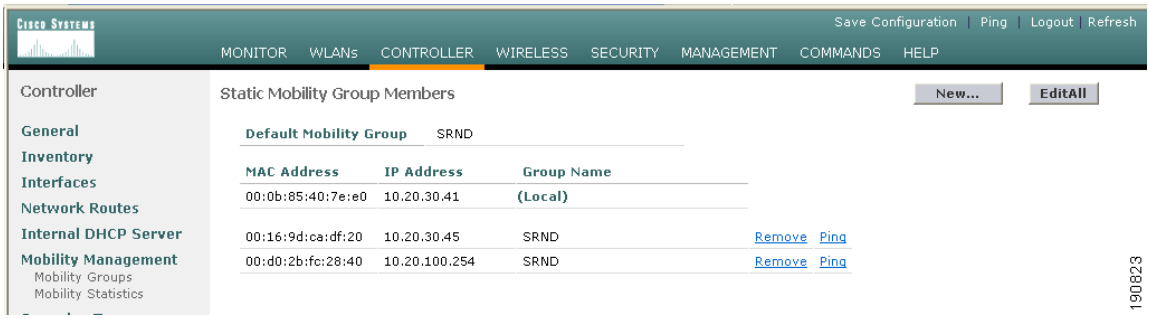

Adding Remote Controllers as Mobility Group Members

Step 3 Click **New** to define a MAC and IP address for each remote controller that will support a guest access WLAN. (See Figure 12-20.)

Figure 12-20 Adding Remote Controllers

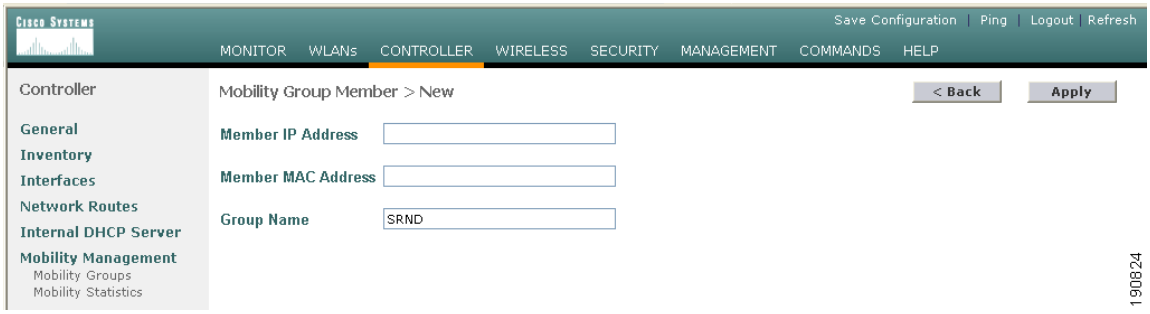

Adding an Anchor Controller as a Mobility Group Member in the Remote Controller

As described in WLAN Anchors and Ethernet in IP to Support Guest Access, page 12-7, each remote controller maps the guest WLAN into an EoIP tunnel that terminates on the anchor controller. Therefore, the anchor controller must be added as a member of the mobility group in each remote controller.

Step 1 Click **New** to add the anchor controller IP and MAC address to the mobility members table.

Step 2 Repeat these steps for each remote controller. (See Figure 12-21.)

Figure 12-21 Adding an Anchor Controller

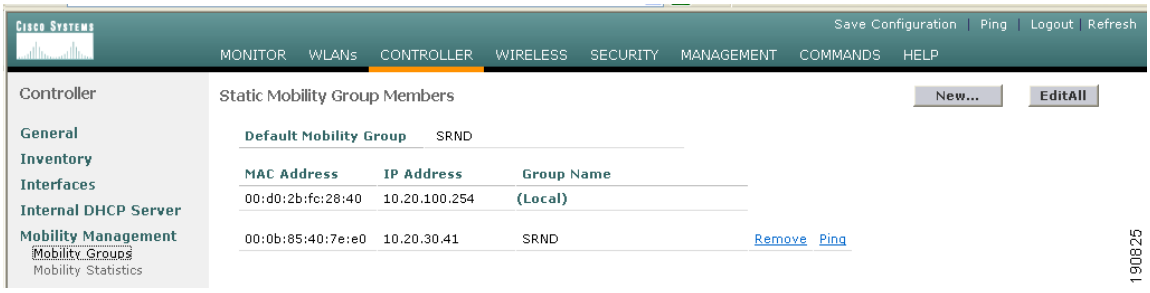

Guest WLAN Configuration

The following section describes how to configure a single guest WLAN. The guest WLAN is configured on each remote controller that manages APs where guest access is required.

Even though the anchor controller is not specifically used to manage lightweight APs, it must also be configured with the guest WLAN because the anchor controller represents an extension of the guest WLAN where user traffic is bridged (using LWAPP between the AP and the remote controller and EoIP between the remote controller and the anchor controller) to a physical interface and VLAN.

Note The parameters defined in the WLAN policies page must be configured identically between the anchor and any given remote controller.

Guest WLAN Configuration for the Remote Controller

Step 1 Click the **WLANs** tab and then click **New**. (See Figure 12-22.)

Figure 12-22 Guest WLAN Configuration

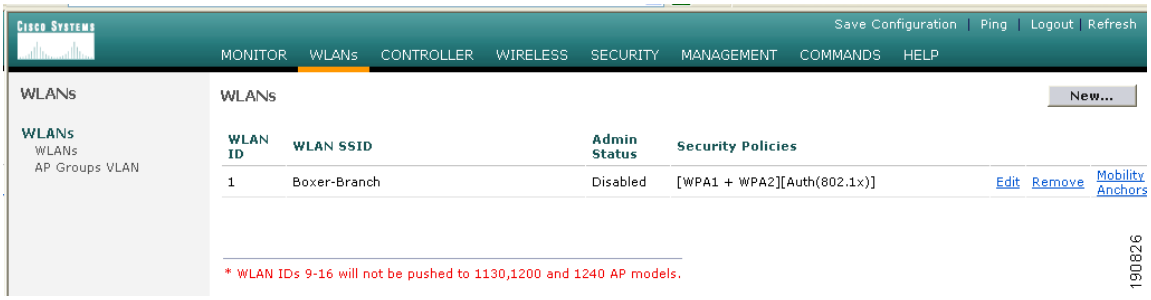

Defining a Guest WLAN SSID

Step 2 Define an SSID that is intuitive or easily recognized by potential guest users.

The controller automatically assigns a VLAN ID. Administrators have the option selecting 1 – 16, as long as the ID is not already in use by another SSID/ WLAN.

Step 3 Click **Apply**. (See Figure 12-23.)

Figure 12-23 Defining a Guest WLAN SSID

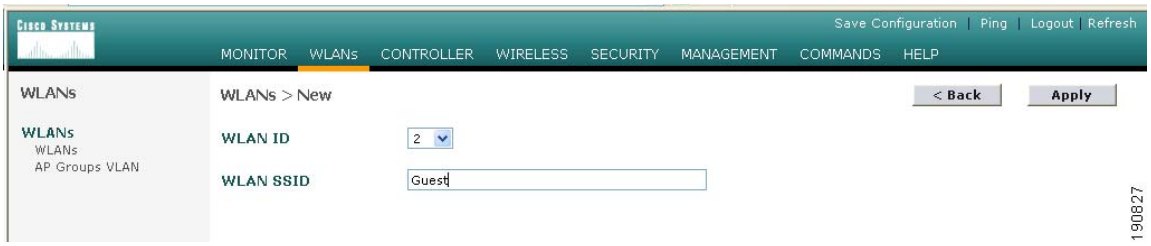

Defining Guest WLAN Policies

The following list represents typical parameter settings used for a guest WLAN:

Step 1 Set the Layer 2 security policy to **None** (802.1x is default).

Step 2 Place a check mark in the Web Policy check box and in either the Authentication check box or Pass-through check box. See Guest User Authentication, page 12-13.

- Broadcast SSID is enabled by default.
- Session time-out default = **0**. (no time out).

  Anything greater than 0 forces de-authentication after expiration and requires the user to re-authenticate through web portal.

- No specific RADIUS servers defined. Authentication is handled by the anchor controller.
- QOS: **Best Effort**.

  This value can be changed.
- DHCP Addr Assignment: **Required**.

  This forces client to use DHCP. Statically configured clients are not accepted.
- Interface name: **Management**.

Note The interface name must be set to the management interface. This causes the remote controller to initiate an EoIP tunnel via its management IP.

Step 3 Click **Apply**. (See Figure 12-24.)

Figure 12-24 Defining Guest WLAN Policies

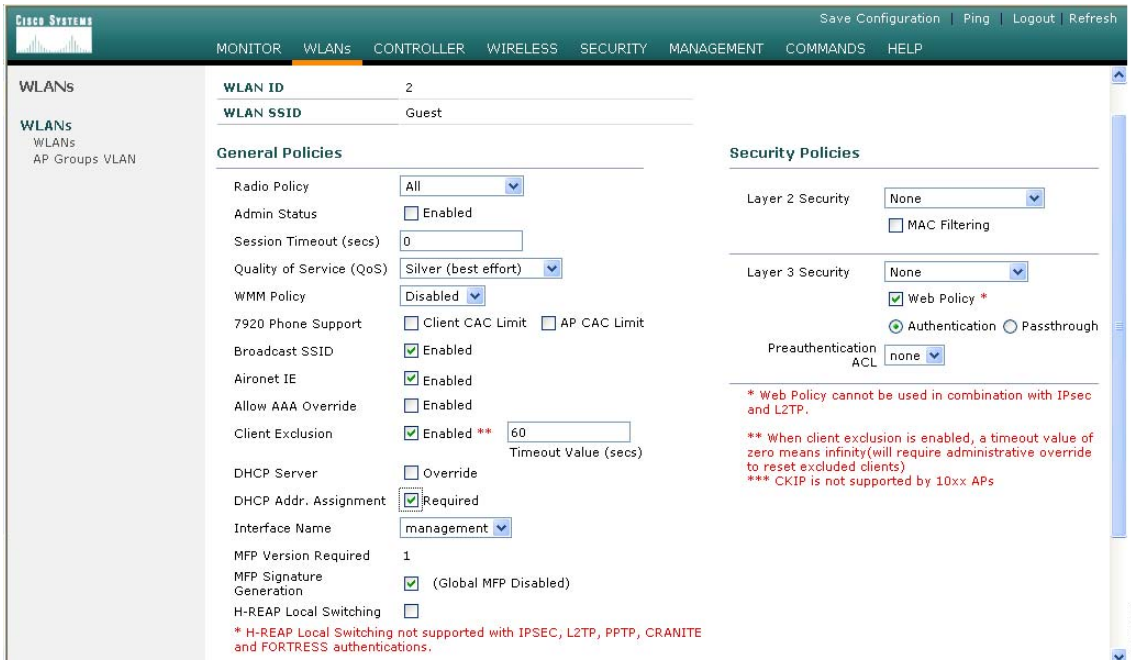

Defining the Guest WLAN Anchor

Step 1 From the WLAN menu, find the newly created guest WLAN.

Step 2 Click **Mobility Anchors**. (See Figure 12-25.)

Figure 12-25 Defining Guest WLAN Anchors

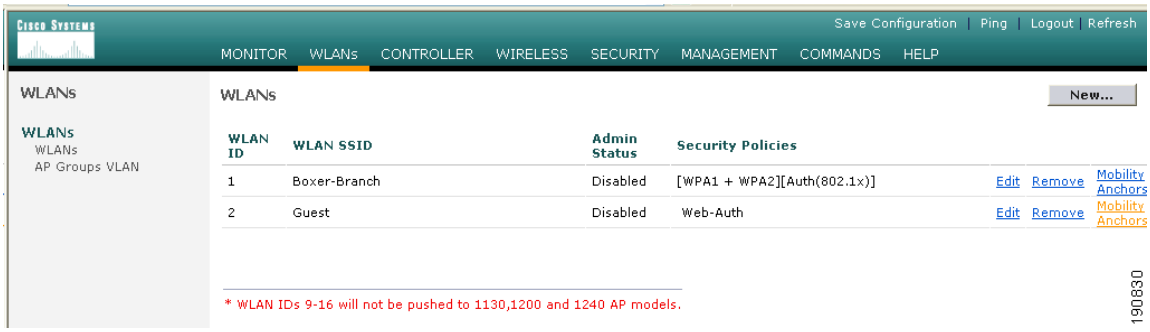

A drop-down list of eligible controller IP addresses is displayed.

Selecting and Creating an Anchor

Step 3 Select the IP address representing the guest access anchor controller deployed in the network DMZ. This is the IP address configured in Adding an Anchor Controller as a Mobility Group Member in the Remote Controller, page 12-23.

Step 4 Click **Mobility Anchor Create**. (See Figure 12-26.)

Figure 12-26 Selecting and Creating an Anchor

Verifying the Guest WLAN Mobility Anchor

The following screenshot shows a mobility anchor assigned to the guest WLAN. You can verify that it can be reached by clicking the **Ping** link as shown above.

Step 5 When finished, click **Back**. (See Figure 12-27.)

Figure 12-27 Verifying the Guest WLAN Mobility Anchor

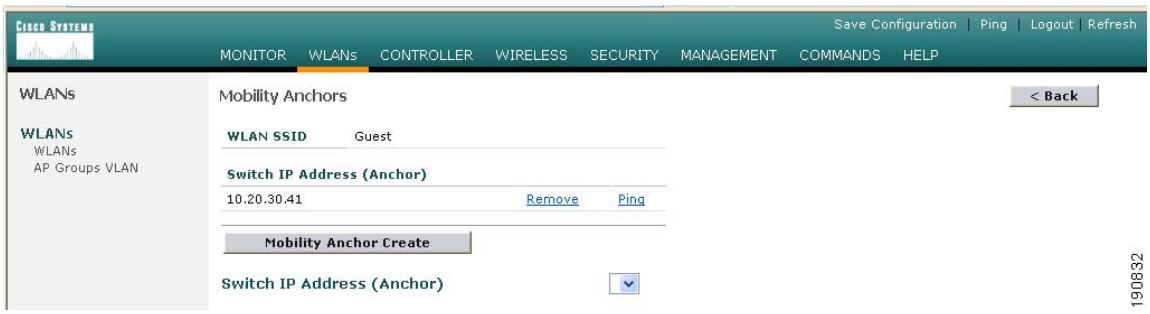

Step 6 After defining the mobility anchor, remember to go back and enable the WLAN.

Enabling the Guest WLAN

Perform the following steps to enable the WLAN.

Step 1 Go back to the WLAN configuration page, find the guest WLAN created in Guest WLAN Configuration for the Remote Controller, page 12-24.

Step 2 Click **Edit**.

Step 3 Click the **Admin Status – Enabled** check box.

Step 4 Click **Apply**. (See Figure 12-28.)

Figure 12-28 Enabling the Guest WLAN

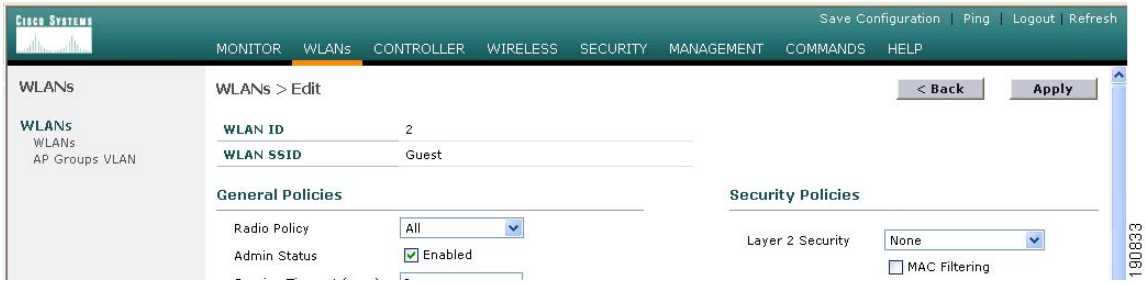

This completes the guest WLAN configuration. Repeat all steps from Guest WLAN Configuration for the Remote Controller, page 12-24 through Enabling the Guest WLAN, page 12-27 for any additional remote controllers.

Guest WLAN Configuration on the Anchor Controller

Guest WLAN configuration on the anchor controller is identical to that performed on the remote controller except for minor differences in the WLAN policies configuration and mobility anchor definition. Repeat all steps from Guest WLAN Configuration for the Remote Controller, page 12-24 through Enabling the Guest WLAN, page 12-27 on the anchor controller.

Note The SSID used for the guest WLAN must be exactly the same as what was configured on the remote controllers.

Guest WLAN Policies for the Anchor Controller

The policies defined for the guest WLAN on the anchor controller are the same except for the interface to which the WLAN is mapped. In this case, the guest WLAN is going to use the interface created in Guest VLAN Interface Configuration, page 12-17.

Step 1 Click the **WLANs** tab.

Step 2 Find the guest WLAN and click **Edit**.

Step 3 Configure the same settings used in Defining a Guest WLAN SSID, page 12-24 except under **Interface Name** choose the interface name created in Guest VLAN Interface Configuration, page 12-17.

Note No RADIUS servers are selected for this WLAN. User authentication of user credentials is handled locally on the controller. See Guest User Authentication, page 12-13.

Step 4 Click **Apply**. (See Figure 12-29.)

Figure 12-29 Configuring Guest WLAN Policies

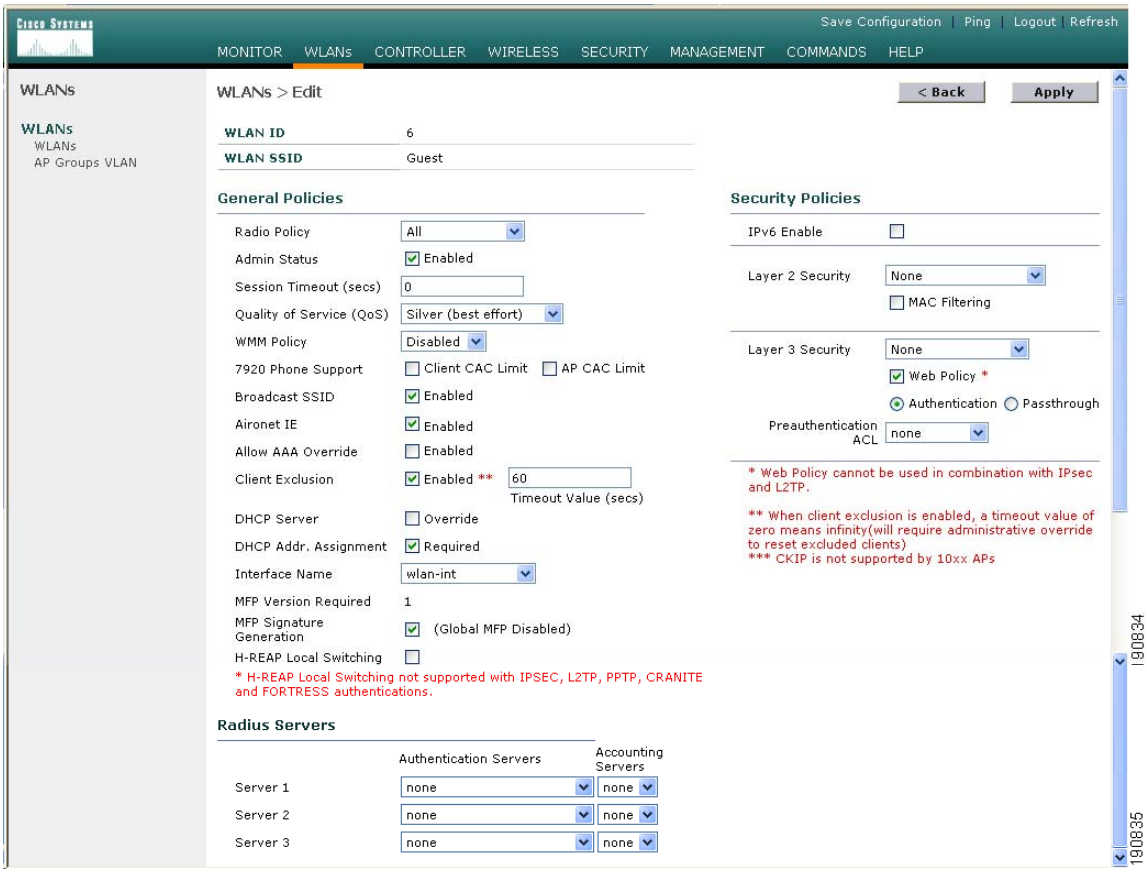

Defining the Guest WLAN Mobility Anchor

The mobility anchor that is used for the guest WLAN is the anchor controller itself.

Step 1 Click the **WLANs** tab.

Step 2 Find the Guest WLAN and click **Mobility Anchors**.

Step 3 From the drop-down list, choose the IP address representing the anchor controller. The IP address has (Local) next to it.

Step 4 Click **Mobility Anchor Create**. (See Figure 12-30.)

Figure 12-30 Defining the Guest WLAN Mobility Anchor

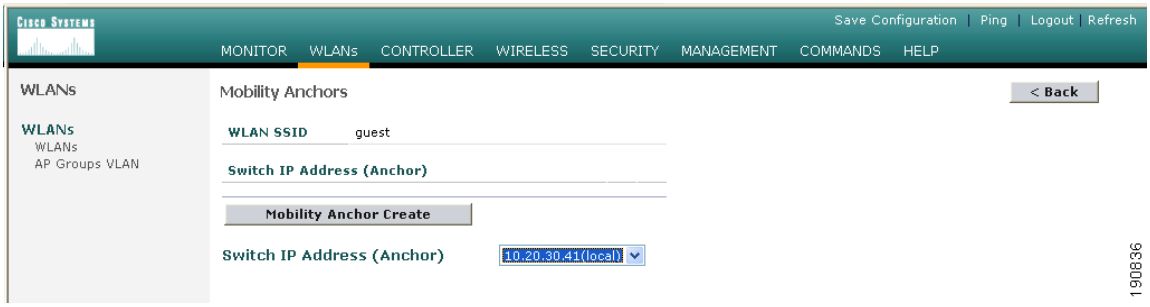

Note that the guest WLAN mobility anchor is *local*. (See Figure 12-31.)

Figure 12-31 Verifying Guest Mobility Anchor is **local**

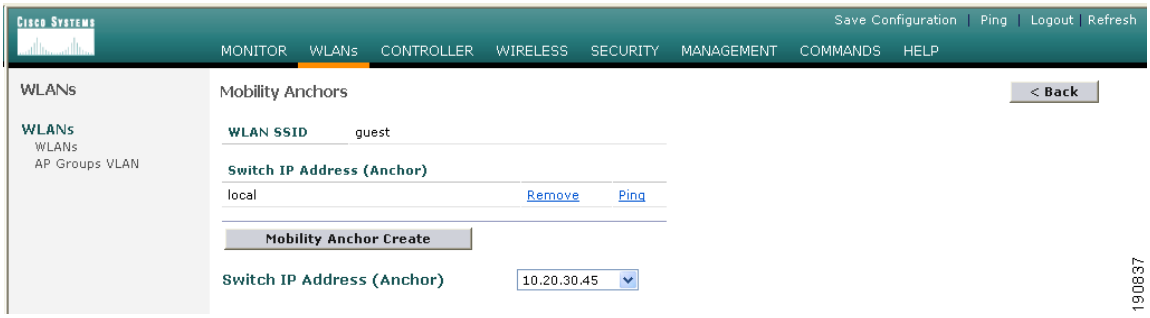

Step 5 Enable the WLAN. Follow the instructions given in Enabling the Guest WLAN, page 12-27.

Web Portal Page Configuration and Management

The internal web server and associated functionality is hosted locally on the anchor controller. When a WLAN is configured to use a web policy, either for authentication or pass-through, the internal web server is invoked by default. No further configuration is required. The internal portal includes a few optional configuration parameters.

Internal Web Page Management

Step 1 Click the **Security** tab.

Step 2 In the left pane, click **Web Login Page**.

The configuration screen shown in Figure 12-32 is displayed. You can change the heading and message information that will appear on the portal page. You can also choose a post authentication redirect URL.

Figure 12-32 Configuration Screen

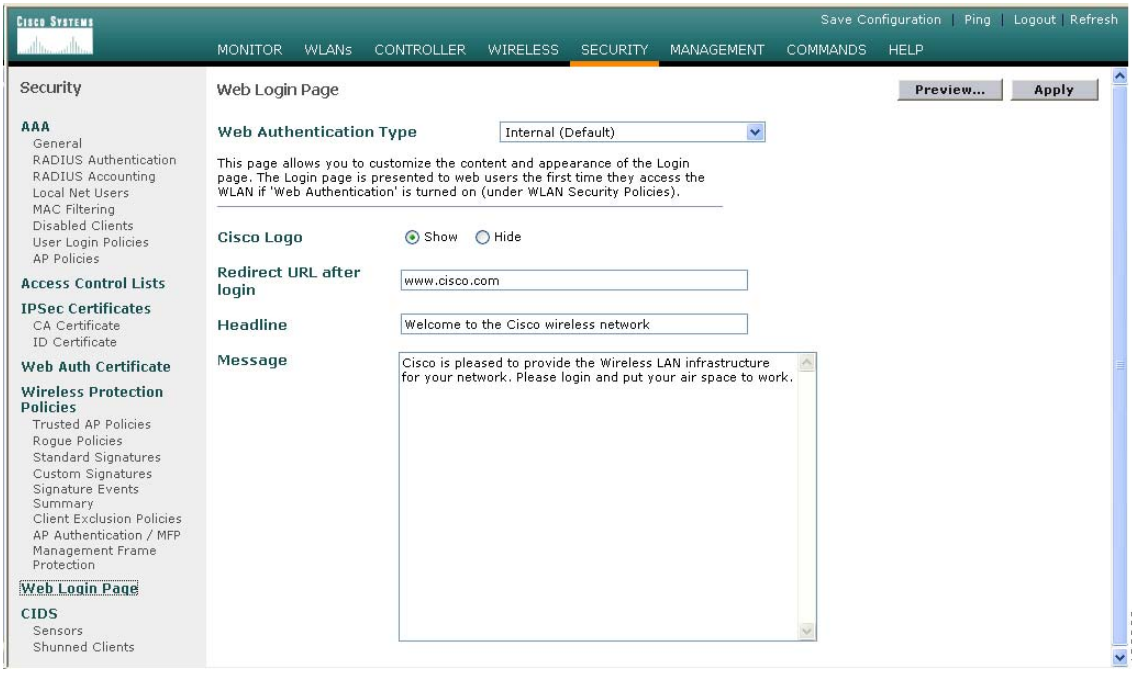

Step 3 Click **Apply**.

Step 4 Optionally, click **Preview** to view what the user will see when redirected.

Importing A Web Page

If you want a customized web page, one can be downloaded and stored locally on the anchor controller. To import a customized web page, perform the following steps:

Step 1 Click the **Commands** tab. (See Figure 12-33.)

Figure 12-33 Importing a Web Page

Step 2 Under File Type choose **Web Auth Bundle**.

Step 3 Define IP address and file path on TFTP server where the files reside.

Step 4 Click **Download** to begin.

There are some caveats to be aware of with regard to downloading a web auth bundle:

- Choose **Web Auth Bundle** from the drop-down list to ensure that the files are stored in the correct directory on the controller.
- The Web Auth Bundle, must be a .tar file of the html and image files being used to create the web login page. When downloaded, the controller will untar the files and place them in the appropriate directory.
- The Web Auth Bundle (tar file) cannot be larger than 1 MB.
- The file name for the html logon page must be **login.html**.

See the following URL for more information about downloading and using customized web pages: http://www.cisco.com/en/US/products/ps6366/products_configuration_guide_book09186a00806b0077.html

Selecting an Imported Web Auth Page

To use a customized web auth page that has been downloaded to the controller, perform the following:

Step 1 Click the **Security** tab.

Step 2 In the left pane, click **Web Login Page**.

Step 3 From the Web Authentication Type drop-down list choose **Customized** (Downloaded)

Step 4 Click **Preview** to view the downloaded page.

Step 5 Click **Apply** when finished. (See Figure 12-34.)

Figure 12-34 Selecting an Imported Web Auth Page

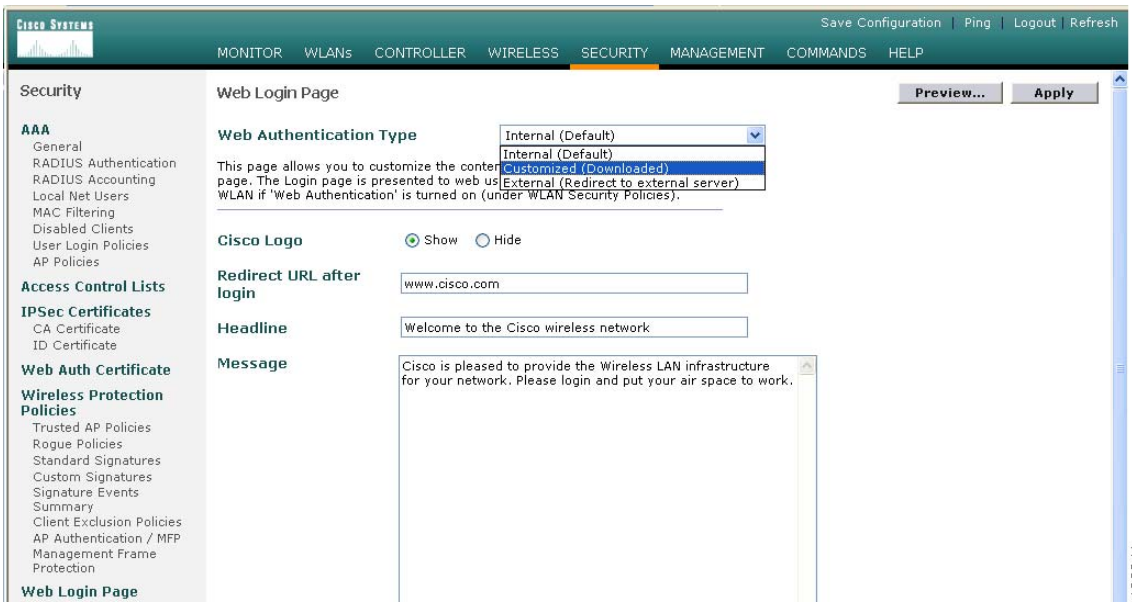

Internal Web Certificate Management

The web auth login page uses SSL for safeguarding user credentials. For simplicity, the controller uses a self-signed certificate. Because the certificate is self-signed, guest users can expect to see a pop-up alert similar to the following when they are redirected to the authentication page shown in Figure 12-35.

Figure 12-35 Authentication Page

At this point, the user can proceed by either clicking **Yes** or they can choose **View Certificate** and manually install it as a trusted site.

The web server uses the virtual interface IP address configured inAnchor Controller Interface Configuration, page 12-17 as its source address. If a host name is defined along with the IP address, that host name must be resolvable by DNS so that:

- The client is redirected to the web auth page.
- The user does not encounter a web certificate error because of conflicts between hostname and host IP address.

Importing an External Web Certificate

In those cases where a legitimate web certificate issued by a trusted root CA is required, one can be downloaded to the controller by performing the following:

Step 1 Click the **Security** tab.

Step 2 In the left pane, click **Web Auth Certificate**. (See Figure 12-36.)

Figure 12-36 Importing an External Web Certificate

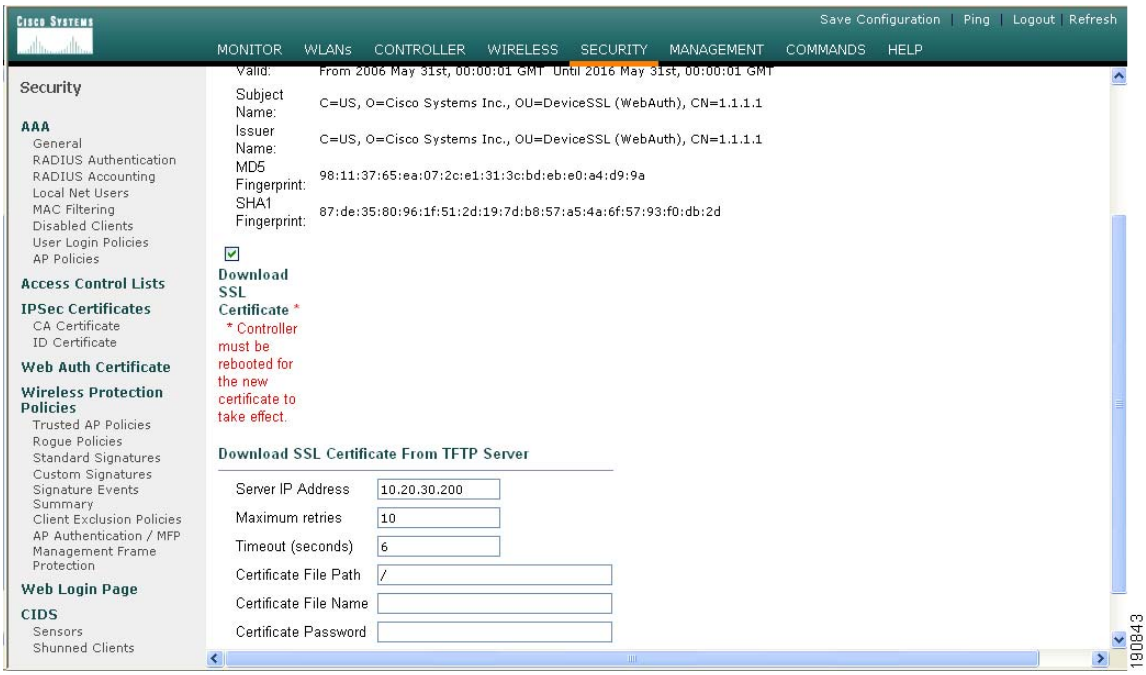

Step 3 Place a check mark in the **Download SSL Certificate** check box

Step 4 Complete the required fields for downloading the certificate.

Step 5 Click **Apply**.

Step 6 After the certificate has been downloaded, reboot the server.

Support for External Web Redirection

In some cases, an enterprise might already have deployed a web portal system to support wired guest access or NAC functionality. If this is the case, the anchor controller can be configured to redirect wireless guest users to an external web portal:

Step 1 Click the **Security** tab.

Step 2 In the left pane, click **Web Login Page**. (See Figure 12-37.)

Figure 12-37 Supporting External Web Redirection

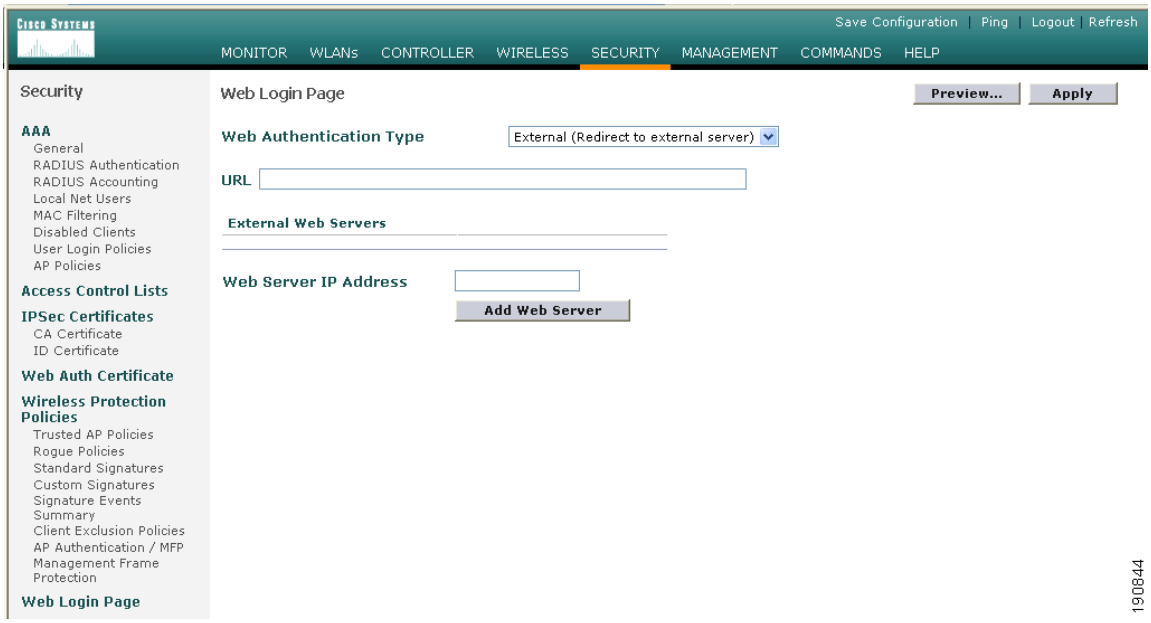

Step 3 Fill in the **Web Server IP** and **URL** fields.

Step 4 Click **Apply**.

See the following URL for more information on the use of external web servers with controller web authentication:
http://www.cisco.com/en/US/products/ps6366/products_configuration_guide_book09186a00806b0077.html

Guest Management

If guest credentials are going to be stored locally on the anchor controller, there are two methods by which they can be created and posted to the controller:

- Through the WCS lobby administration interface
- Directly on the controller through a local lobby admin account

Guest Management Using WCS

The following configuration examples assume WCS version 4.0 has been installed, configured, and a lobby ambassador account has been established. See Chapter 8, "Cisco Unified Wireless Control System," and the following URL for more information on installing and configuring WCS: http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_book09186a00806b57ec.html

Step 1 Log in to WCS using the Lobby Ambassador credentials assigned by the system administrator. (See Figure 12-38.)

Figure 12-38 Using WCS

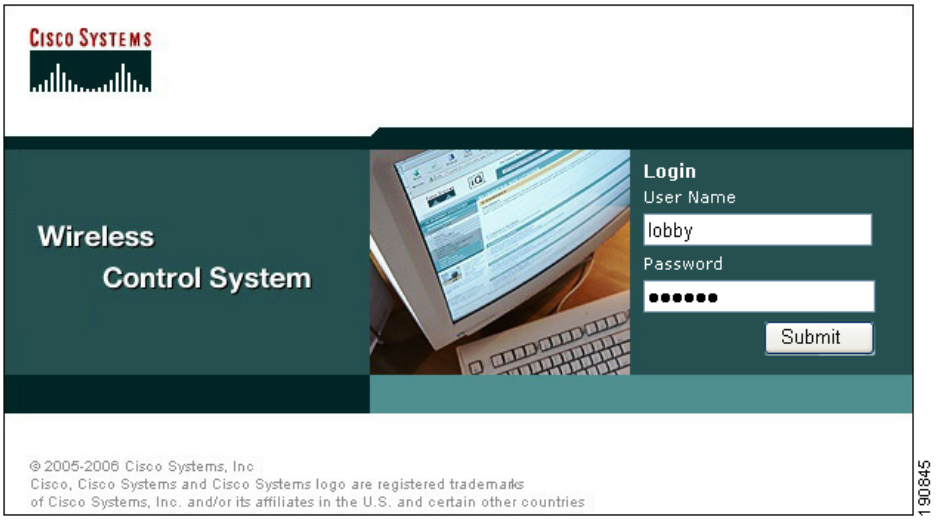

After logging in, the screen shown in Figure 12-39 is displayed.

Figure 12-39 Adding a Template

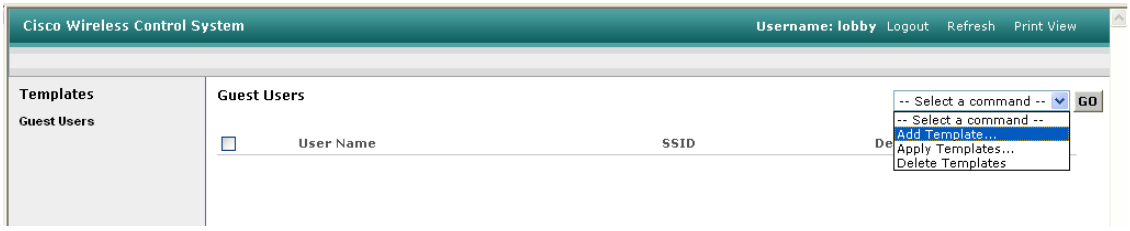

Step 2 To add a user template, from the drop-down list, choose **Add Template** and click **Go**.

The screen shown in Figure 12-40 appears.

Figure 12-40 Guest User Template

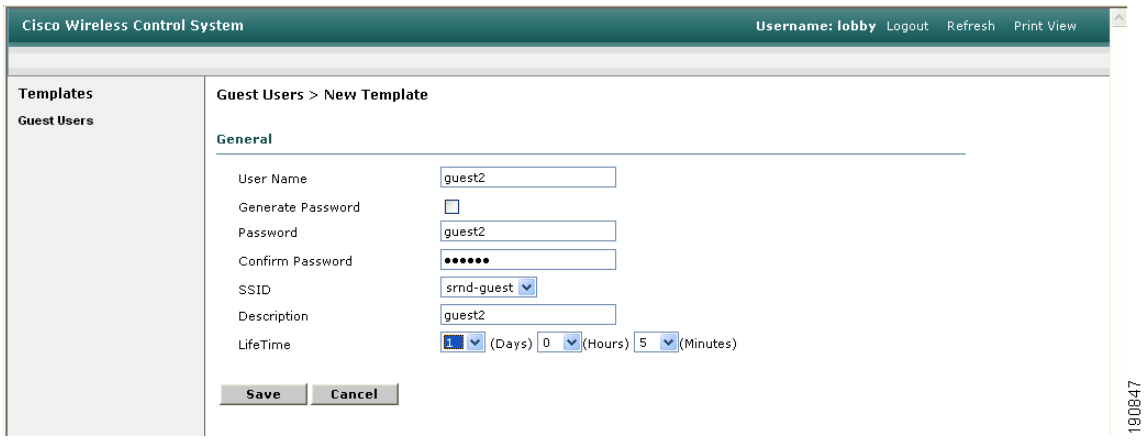

Step 3 To creating user credentials, enter a username and password (auto or manual).

Step 4 Select the WLAN / SSID the guest account applies to. (only WLANs configured to use a web policy are displayed).

Step 5 Enter the lifetime for credentials.

Step 6 Enter a description for the user.

Step 7 Click **Save**.

Applying Credentials

After the credentials have been created, the screen shown in Figure 12-41 offers the option to apply them to one or more controllers.

Figure 12-41 Applying Credentials

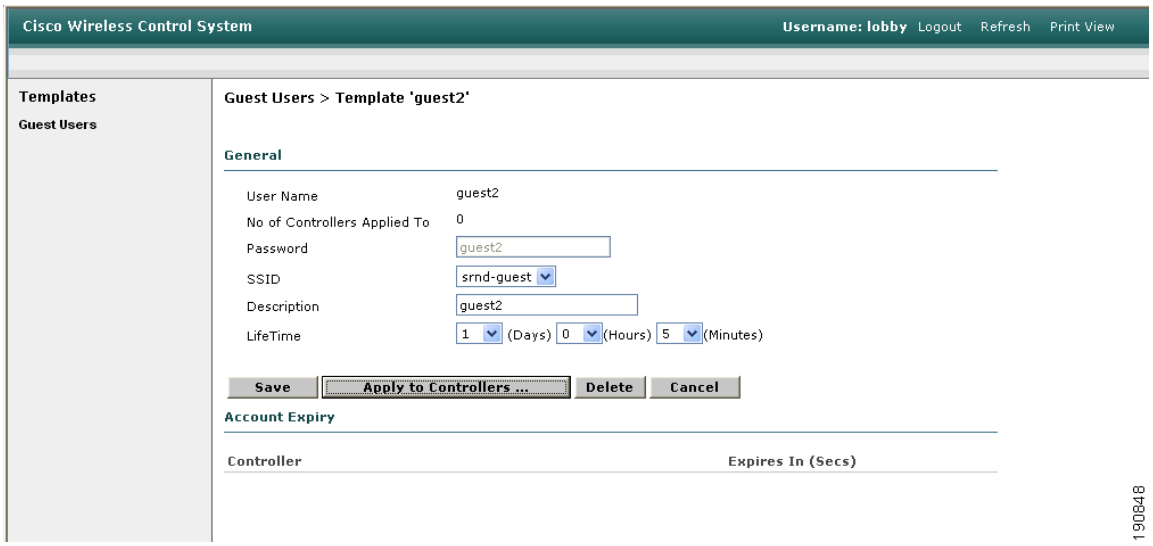

Step 1 Click **Apply to Controllers**.

As shown in Figure 12-42, a list of eligible controllers is displayed (only those controllers that have been configured with the guest WLAN are displayed).

Figure 12-42 Apply to Controllers Screen

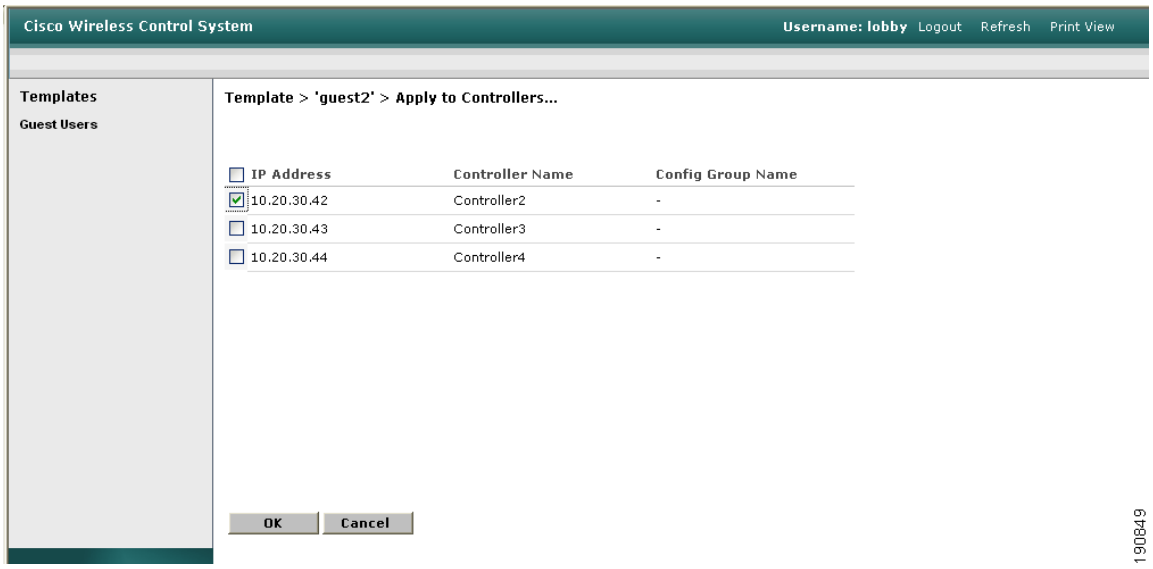

Note The guest WLAN will have been configured on the anchor controller and one or more remote controllers. The anchor controller is the only controller where the guest credentials need to be applied.

Step 2 To apply credentials to the anchor controller, from the list of controllers, choose the anchor controller, and click **OK**.

A confirmation page that verifies that the user credentials were saved to the anchor controller is displayed. (See Figure 12-43.)

Figure 12-43 Confirmation Page

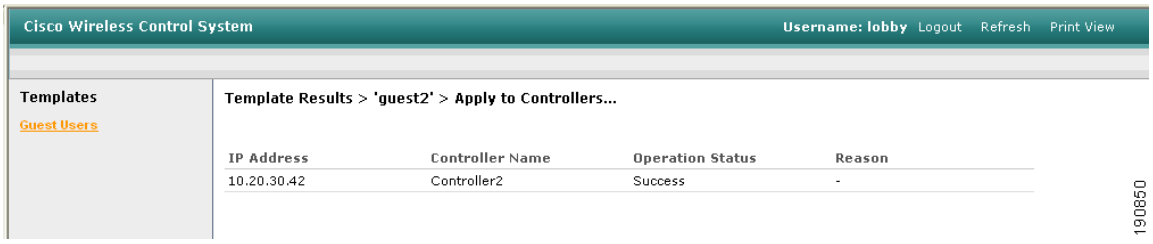

Step 3 In the left pane, click **Guest Users** to return to the summary page. (See Figure 12-44.)

Figure 12-44 Summary Page

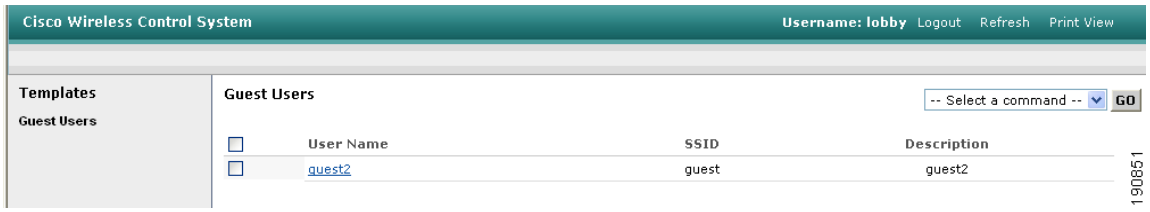

Step 4 To edit guest credentials, from the summary page, click the user name that you want to edit.

The user template is displayed, as shown in Figure 12-45.

Figure 12-45 Editing Guest Credentials

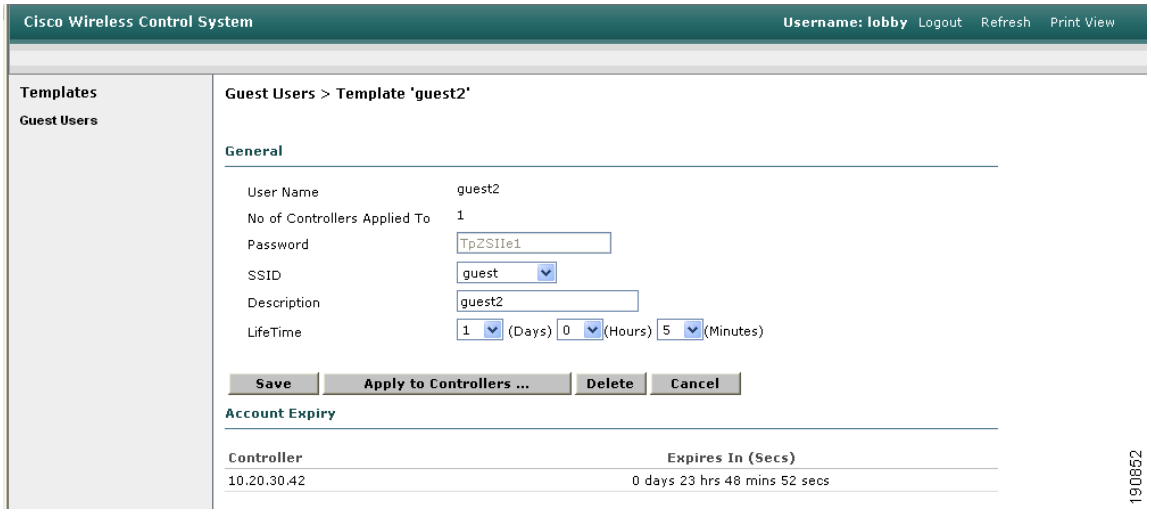

In this page, you can make the following modifications if desired:

- Change the WLAN to which the credentials apply.
- Change the user description.
- Change the Lifetime of the credentials.
- Apply the credentials to other controllers.
- Delete the user template.

Note If a user template is deleted from WCS while a user is active, they are de-authenticated.

Managing Guest Credentials Directly on the Anchor Controller

The following procedure assumes a network administrator has established a Local Management User account with Lobby Admin privileges on the controller.

Step 1 Log in to the anchor controller using the lobby admin credentials assigned by the system administrator. Remember that conduits might need to be opened through a firewall to permit HTTP/HTTPS for web administration of the controller. See Anchor Controller Positioning, page 12-9.

After login, the screen shown in Figure 12-46is displayed.

Figure 12-46 Anchor Controller Login

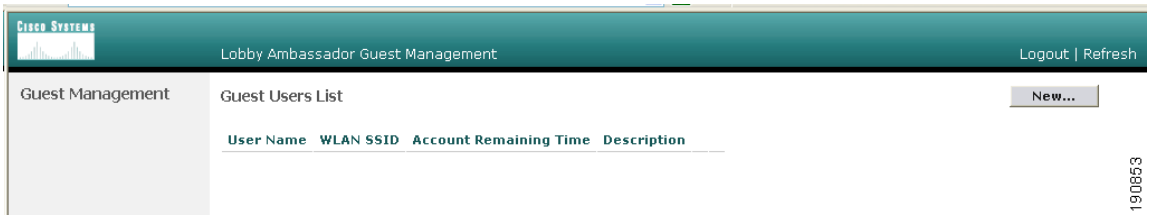

Step 2 Click **New**.

The screen shown in Figure 12-47 appears.

Figure 12-47 Creating User Credentials

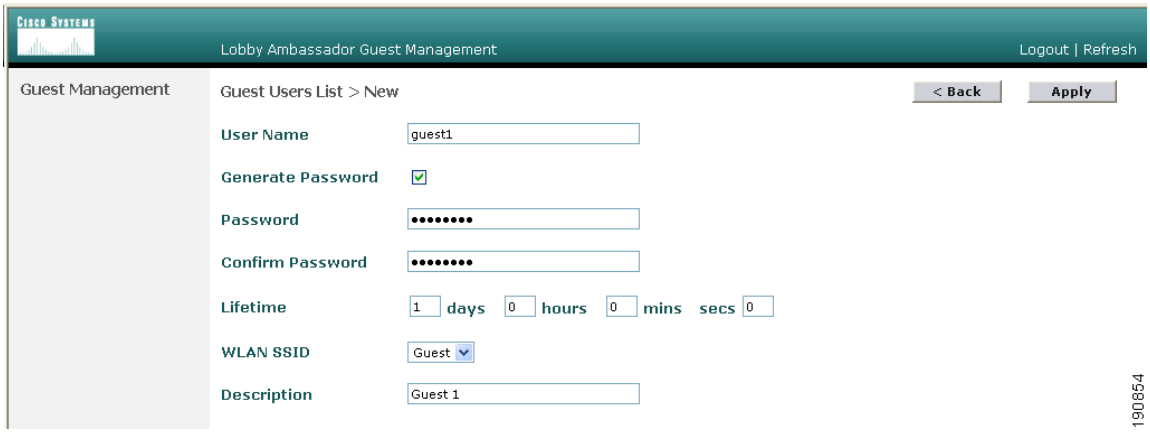

Step 3 To create user credentials, perform the following steps:

a. enter a username and password (manual or auto).

b. Select the WLAN/SSID to which the guest account applies (only WLANs configured with web policy will be displayed).

c. Enter a lifetime for the credentials.

d. Enter a description for the user.

Step 4 Click **Apply**.

The screen shown in Figure 12-48 appears and shows the newly added guest user.

Figure 12-48 Guest Users List

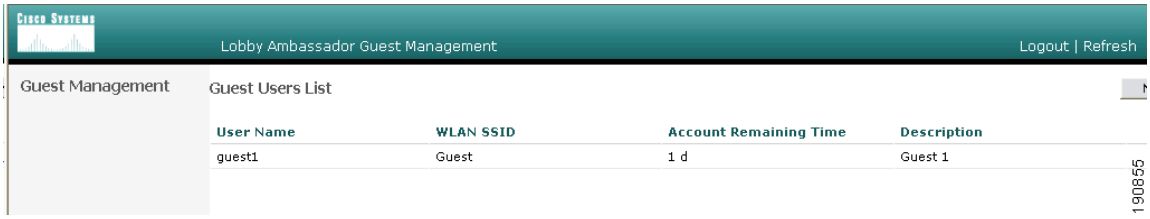

From this screen you have the option to do the following:

- Edit the existing user
- Delete the existing user
- Add a new user

Configuring the Maximum Number of User Accounts

The default number of guest user accounts that can be defined on the controller is 512. This value can be changed by completing the following steps.

Step 1 Click the **Security** tab. (See Figure 12-49.)

Figure 12-49 Configuring the Maximum Number of User Accounts

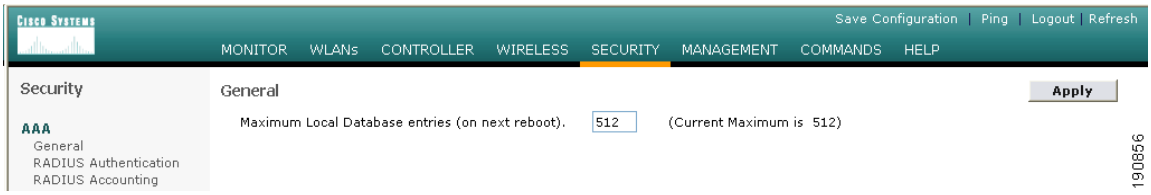

Step 2 In the left pane, click **General** under AAA properties.

Step 3 Configure the maximum number of user database entries (between 512 and 2048).

Step 4 Click **Apply**.

Guest User Management Caveats

- Guest credentials can be added using either method above or both methods together.
- When using WCS, the lobby admin does not have visibility of user accounts that might have been created locally on the anchor controller. If a WCS lobby admin attempts to add a user name that is already in use, a pop-up window prompting the admin to choose a different user name is displayed.
- When adding user accounts locally on the controller, the admin sees all accounts that have been created, including those that were created via WCS.
- If a guest user is currently authenticated to a WLAN and their credentials are deleted from WCS or locally on the controller, the user traffic stops flowing, and the user is de-authenticated.

External Radius Authentication

As described in Guest User Authentication, page 12-13, an external RADIUS server can be used to authenticate guest users in place of creating and storing guest credentials locally on the anchor controller. If this method is used, then the lobby admin features described in Guest Management, page 12-35 cannot be used. It is assumed that some other guest management system will be used in conjunction with the external RADIUS server.

To configure a guest WLAN to use an external RADIUS server, perform the following configuration steps on the anchor controller.

Adding a RADIUS Server

Step 1 Click the **Security** tab.

A summary screen is displayed. (See Figure 12-50.)

Figure 12-50 Summary Screen

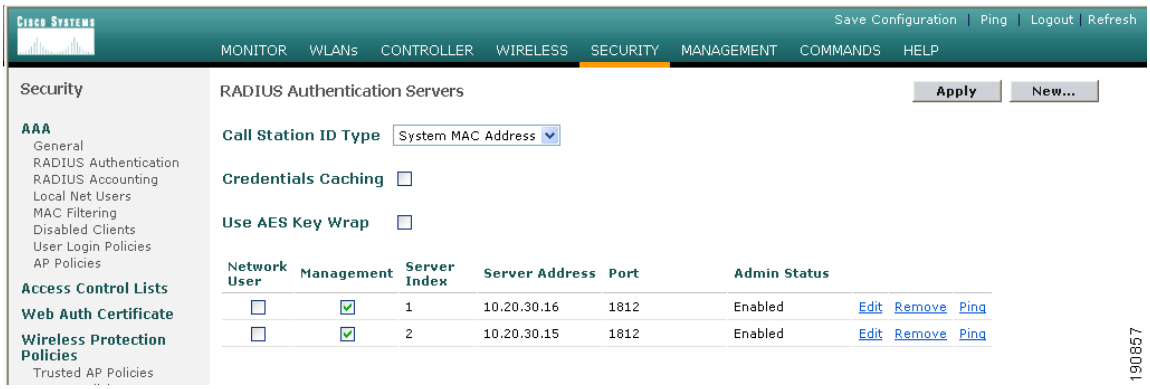

Step 2 Click **New**.

The screen shown in Figure 12-51appears.

Figure 12-51 Defining RADIUS Server Settings

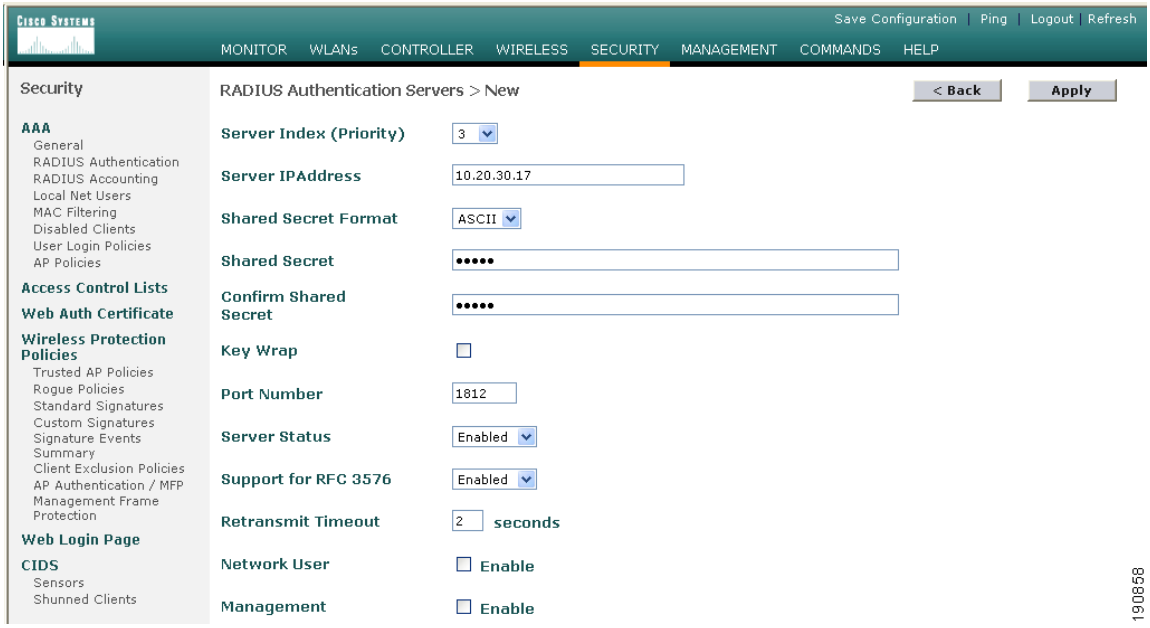

Step 3 To define RADIUS server settings, configure the IP address, Shared Secret, and Authentication Port Number as defined on the RADIUS server.

If the Network User check box is cleared, the RADIUS server is used only for user authentication when it is specifically selected under a given WLANs RADIUS setting. Otherwise, if the check box is checked, the server is used globally for all user authentications based on its server priority.

Step 4 Click **Apply**.

The summary screen shows the newly added server. (See Figure 12-52.)

Figure 12-52 Summary Screen

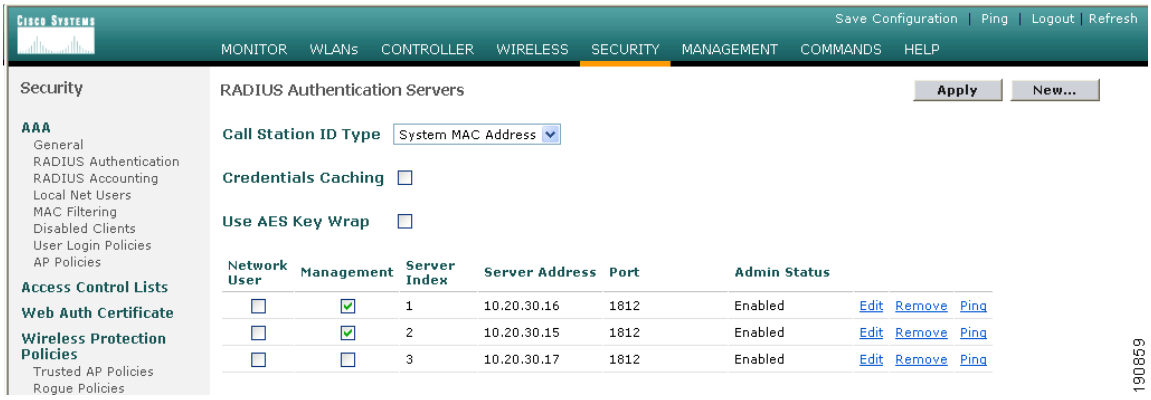

Step 5 To select a RADIUS server, click the **WLANs** tab.

Step 6 The screen shown in Figure 12-53 appears.

Figure 12-53 WLANs Tab

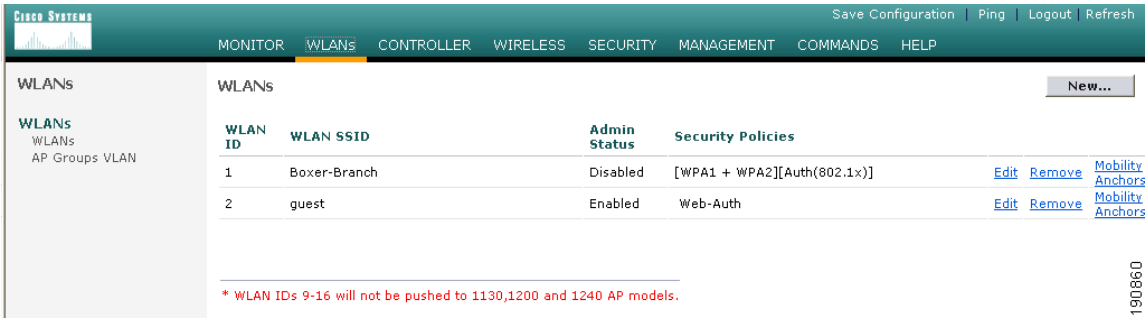

Step 7 Find the guest WLAN and click **Edit**.

The guest WLAN configuration screen is displayed, as shown in Figure 12-54.

Figure 12-54 Guest WLAN Configuration Screen

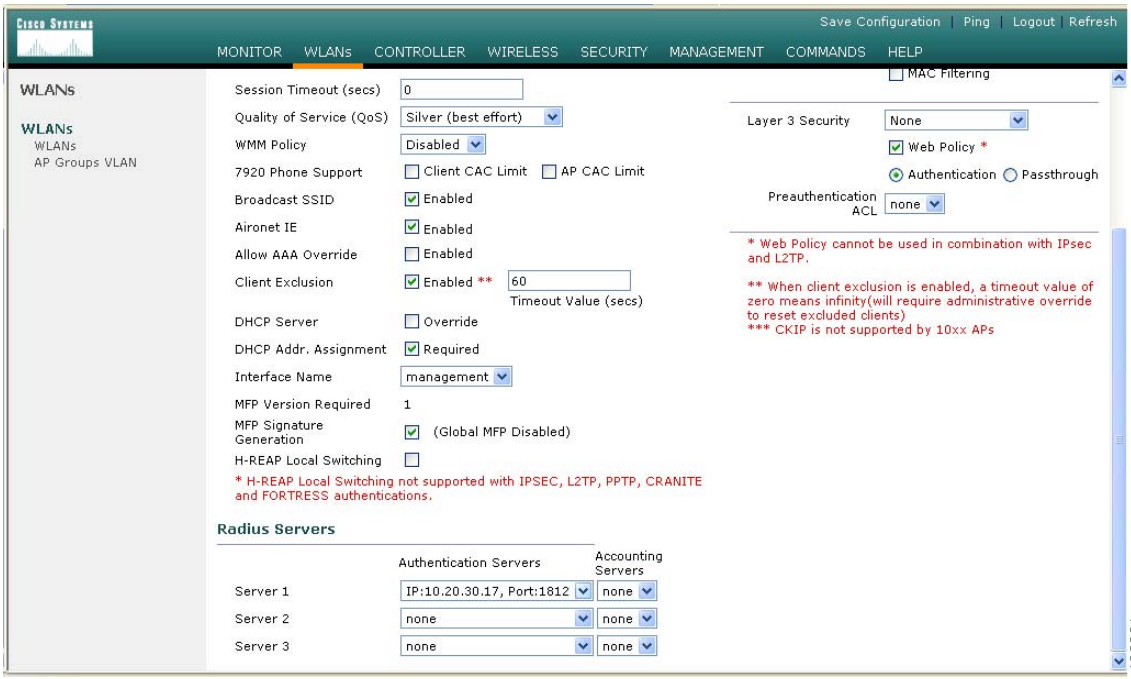

Step 8 Select the RADIUS server to be used for web authentication from the drop-down list.

External Access Control

The centralized guest access topology described in this chapter can be integrated with an external access control platform such as BBSM or Clean Access.

In this scenario, an enterprise might have already deployed an access control platform in their Internet DMZ to support wired guest access services (see Figure 12-55).

Figure 12-55 Wireless Guest Access with External Access Control

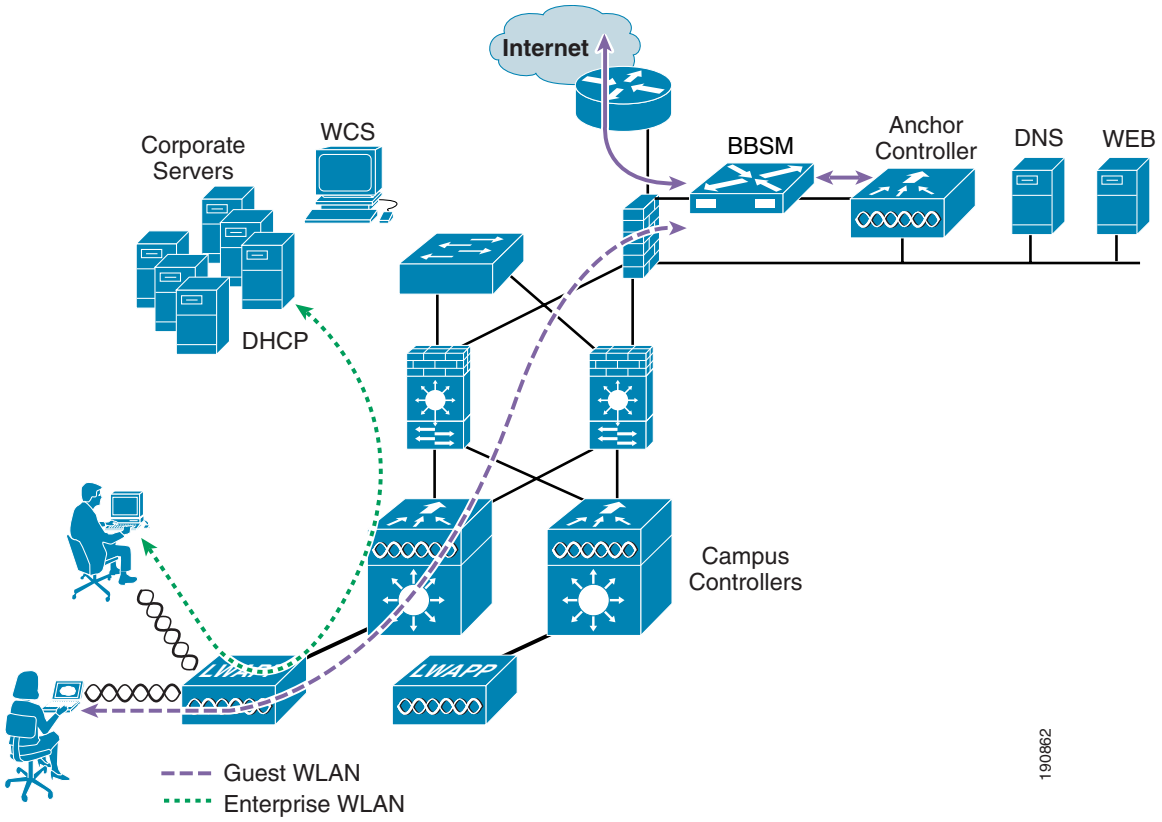

As shown in Figure 12-55, The wireless guest access topology remains the same except the guest VLAN interface on the anchor controller, instead of connecting to a firewall or border router, connects to an inside interface on an access control platform such as BBSM.

In this scenario, the BBSM platform is responsible for redirection, web authentication and subsequent access to the Internet. The campus and anchor controllers are only used to tunnel guest WLAN traffic across the enterprise into the DMZ where BBSM or some other platform is used to manage guest access.

Configuration of the guest WLAN, campus and anchor controllers is the same as described in the previous examples. The only exception is that web policy is disabled under the guest WLANs security policy (see Figure 12-56).

Figure 12-56 Guest WLAN Security Policy

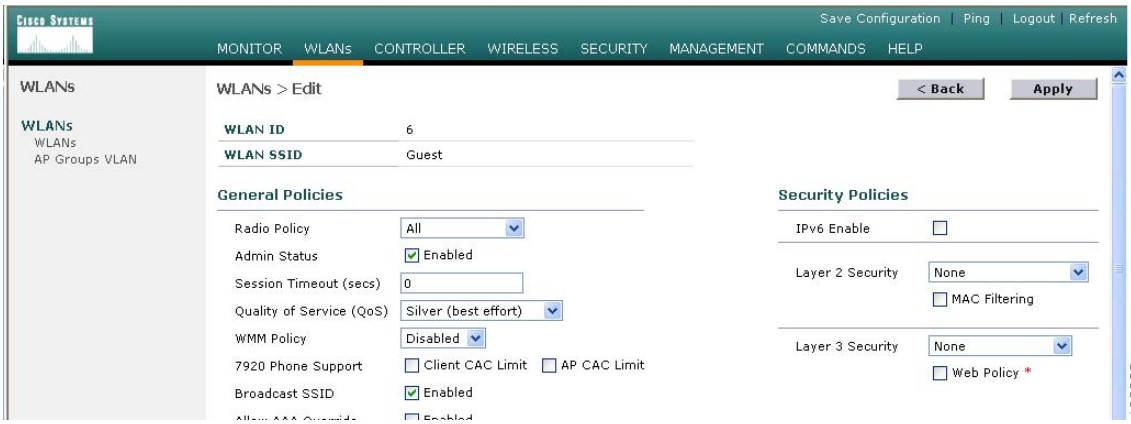

The configuration above establishes a WLAN with no security policies. Guest traffic passes through the anchor controller to the inside or untrusted interface of BBSM or Cisco Clean Access Server, where it is blocked until a user has authenticated.

DHCP can be hosted locally on the controller or externally via BBSM or dedicated server.

Its beyond the scope of this chapter to address BBSM or other external platform specific configuration. See the platform documentation for additional configuration guidelines.

Verifying Guest Access Functionality

The guest access service is working correctly if a user:

- Can associate to the guest WLAN
- Receives an IP address via DHCP
- Opens their browser and is redirected to the web authentication page
- Enters their credentials and connects to the Internet (or other upstream services)

Troubleshooting Guest Access

The following verifications and troubleshooting tasks assume:

- The solution is using the web authentication functionality resident in the anchor controller
- User credentials are created and stored locally on the anchor controller

Before attempting to troubleshoot the various symptoms below, at the very least you should be able to ping from the campus (foreign) controller to the anchor controller. If not, verify routing.

Next, you should be able to perform the following advanced pings. These can only be performed via the serial console interfaces of the controllers:

- **mping** *neighbor WLC ip*

  This pings the neighbor controller through the LWAPP control channel.

- **eping** *neighbor WLC ip*

  This pings the neighbor controller through the LWAPP data channel.

If pings go through, but mpings do not, ensure that each WLCs mobility group name is the same and ensure that each WLCs IP, MAC, and mobility group name is entered in every WLC mobility list.

If pings and mpings are successful, but epings are not, check the network to make sure that IP protocol 97 (Ethernet-over-IP) is not being blocked.

User Cannot Associate to the Guest WLAN

- Verify the guest WLAN is enabled on the anchor controller and all remote controllers that support the guest WLAN
- Verify the guest WLAN SSID is being broadcast.
- Verify client adapter/software configuration

User Does Not Obtain an IP Address via DHCP

- Verify WLAN configuration settings are identical on the anchor and remote controllers
- Verify the guest WLAN is enabled on the anchor controller and all remote controllers that support the guest WLAN
- Check for a proper DHCP server address under the guest VLAN interface settings on the anchor controller
  - If using an external DHCP server, the IP address should be that of the external server.
  - Verify reachability to the external DHCP server from the anchor controller
  - If using the anchor controller for DHCP services, the DHCP server IP address should be the controllers management IP address.
  - Verify that a DHCP scope has been configured and enabled on the controller
  - Verify that the DHCP scopes network mask is consistent with the mask on the guest VLAN interface.
  - Verify the DHCP scope does not overlap with any addresses assigned to network infrastructure

User is Not Redirected to Web Auth Page

The following assumes the user is able to associate to the guest WLAN and obtain an IP address

- Verify valid DNS servers are being assigned to the client via DHCP
- Ensure the DNS servers are reachable from the anchor controller
- Verify the URL being opened in the web browser is resolvable
- Verify the URL being opened in the web browser is connecting to http port 80

Note The internal web auth server does not redirect incoming requests on ports other than 80.

User Cannot Authenticate

- Verify user credentials are active on the anchor controller.
  - Guest credentials typically have a lifetime associated with them. If the credentials have expired, they do not appear under the Local Net Users list on the anchor controller. Use WCS to re-apply the user template or re-create user credentials locally on the controller. See Guest Management Using WCS, page 12-36 and Guest Credentials Management, page 12-12.

- Verify user password

User Cannot Connect to Internet or Upstream Service

- Verify routing to and from the anchor controller from the firewall or border router connecting to the controller
- Verify NAT configuration on firewall or Internet border router (if applicable)

System Monitoring

Following are some monitoring commands that might be helpful in troubleshooting.

Anchor Controller

From the serial console port:

```
(Cisco Controller) >show client summary

Number of Clients............................... 1
MAC Address        AP Name            Status         WLAN  Auth  Protocol  Port
-----------------  -----------------  -------------  ----  ----  --------  ----
00:40:96:a6:d5:3a  10.20.100.254      Associated     6     Yes   Mobile    1
```

Note that the protocol is Mobile. The Auth field reflects the actual status of the user. If the user has passed web auth, the field displays YES. If not, the field shows NO.

Also notice the AP name. This is the management IP address of the remote controller (originating controller).

From the summary information, use the client MAC to show more detail:

```
(Cisco Controller) >show client detail 00:40:96:a6:d5:3a

Client MAC Address.............................. 00:40:96:a6:d5:3a
Client Username ................................ guest1
AP MAC Address.................................. 00:00:00:00:00:00
Client State.................................... Associated
Wireless LAN Id................................. 6
BSSID........................................... 00:00:00:00:00:05
Channel......................................... N/A
IP Address...................................... 10.20.31.101
Association Id.................................. 0
Authentication Algorithm........................ Open System
Reason Code..................................... 0
Status Code..................................... 0
Session Timeout................................. 0
Client CCX version.............................. No CCX support
Re-Authentication Timeout....................... 81586
Remaining Re-Authentication Time................ 79010
Mirroring....................................... Disabled
QoS Level....................................... Silver
Diff Serv Code Point (DSCP)..................... disabled
802.1P Priority Tag............................. disabled
WMM Support..................................... Disabled
Mobility State.................................. Export Anchor
Mobility Foreign IP Address..................... 10.20.100.254  <Remote Controller
Mobility Move Count............................. 1
Security Policy Completed....................... Yes
Policy Manager State............................ RUN
Policy Manager Rule Created..................... Yes
```

```
NPU Fast Fast Notified..... Yes
Policy Type..... N/A
Encryption Cipher..... None
EAP Type..... Unknown
Interface..... wlan-int
VLAN..... 31
Client Capabilities:
      CF Pollable..... Not implemented
      CF Poll Request..... Not implemented
      Short Preamble..... Not implemented
      PBCC..... Not implemented
      Channel Agility..... Not implemented
      Listen Interval..... 0
Client Statistics:
      Number of Bytes Received..... 0
      Number of Bytes Sent..... 0
      Number of Packets Received..... 0
      Number of Packets Sent..... 0
      Number of Policy Errors..... 0
      Radio Signal Strength Indicator..... Unavailable
      Signal to Noise Ratio..... Unavailable
Nearby AP Statistics:
        TxExcessiveRetries: 0
        TxRetries: 0
        RtsSuccessCnt: 0
        RtsFailCnt: 0
        TxFiltered: 0
        TxRateProfile: [0,0,0,0,0,0,0,0,0,0,0,0]
```

The same information can be obtained through the web configuration and management interface of the controller under **Client Detail**. (See Figure 12-57.)

Figure 12-57 Monitor > Client Detail

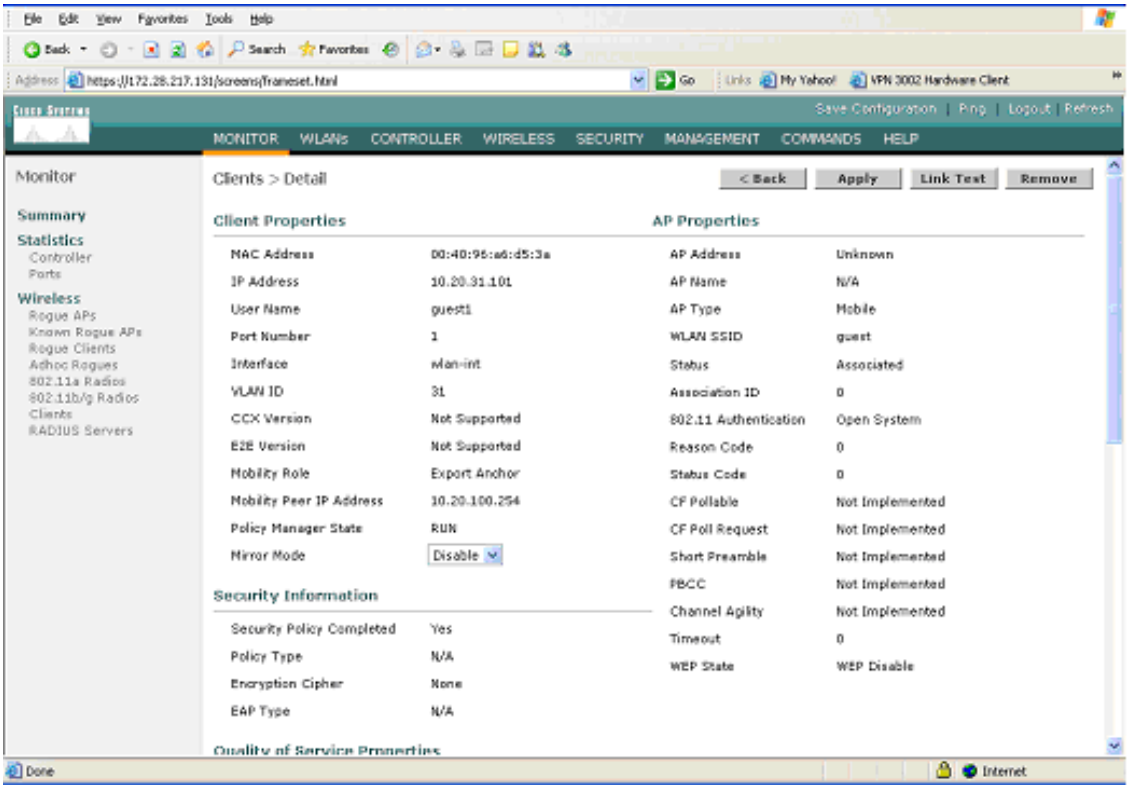

Campus (Foreign) Controller

From the serial console port:

```
(Cisco Controller) >show client summary

Number of Clients.............................. 1
MAC Address        AP Name            Status         WLAN  Auth  Protocol  Port
-----------------  -----------------  -------------  ----  ----  --------  ----
00:40:96:a6:d5:3a  Branch:24:a6:50    Associated     2     Yes   802.11g   1
```

Note that the protocol field is 802.11g, whereas the protocol field on the anchor controller for the same client is mobile. The campus (foreign) controller always shows the user as authenticated and the AP name reflects the actual AP to which the client is associated.

Additional details can be obtained using the following:

```
Cisco Controller) >show client detail 00:40:96:a6:d5:3a

Client MAC Address.............................. 00:40:96:a6:d5:3a
Client Username ................................ N/A
AP MAC Address.................................. 00:0b:85:24:a6:50
Client State.................................... Associated
Wireless LAN Id................................. 2
BSSID........................................... 00:0b:85:24:a6:5e
Channel......................................... 1
IP Address...................................... Unknown
Association Id.................................. 4
Authentication Algorithm........................ Open System
```

```
Reason Code..................................... 0
Status Code..................................... 0
Session Timeout................................. 0
Client CCX version.............................. No CCX support
Re-Authentication Timeout....................... 0
Remaining Re-Authentication Time................ Timer is not running
QoS Level....................................... Silver
Diff Serv Code Point (DSCP)..................... disabled
802.1P Priority Tag............................. disabled
WMM Support..................................... Disabled
Mobility State.................................. Export Foreign
Mobility Anchor IP Address...................... 10.20.30.41<anchor controller>
Mobility Move Count............................. 0
Security Policy Completed....................... Yes
Policy Manager State............................ RUN
Policy Manager Rule Created..................... Yes
Policy Type..................................... N/A
Encryption Cipher............................... None
EAP Type........................................ Unknown
Interface....................................... management <source of EoIP Tunnel>
VLAN............................................ 0
Client Capabilities:
      CF Pollable............................... Not implemented
      CF Poll Request........................... Not implemented
      Short Preamble............................ Implemented
      PBCC...................................... Not implemented
      Channel Agility........................... Not implemented
      Listen Interval........................... 0
Client Statistics:
      Number of Bytes Received.................. 83288
      Number of Bytes Sent...................... 310361
      Number of Packets Received................ 670
      Number of Packets Sent.................... 430
      Number of Policy Errors................... 0
      Radio Signal Strength Indicator........... -30 dBm
      Signal to Noise Ratio..................... 62 dB
Nearby AP Statistics:
        TxExcessiveRetries: 0
        TxRetries: 0
        RtsSuccessCnt: 0
        RtsFailCnt: 0
        TxFiltered: 0
        TxRateProfile: [0,0,0,0,0,0,0,0,0,0,0,0]
      Branch:24:a6:50(slot 1) ...................
antenna0: 451 seconds ago -26 dBm.............. antenna1: 1522 seconds ago -68 dBm
```

The same information can be obtained through the controller web configuration and management interface under client detail (see Figure 12-57).

Debug Commands

Additional debug commands that might be used from the serial console include the following:

```
debug mac addr <client mac address>
debug mobility handoff enable
debug mobility directory enable
debug dhcp packet enable
debug pem state enable
debug pem events enable
debug dot11 mobile enable
debug dot11 state enable
```

CHAPTER 13

Mobile Access Router, Universal Bridge Client, and Cisco Unified Wireless

The Cisco 3200 Series Mobile Access router (also referred to as the MAR3200) is a compact, high-performance access solution that offers seamless mobility and interoperability across wireless networks. The size of the Cisco MAR3200 (see Figure 13-1) makes it ideal for use in vehicles in public safety, homeland security, and transportation sectors. The MAR3200 delivers seamless communications mobility across multiple radio, cellular, satellite, and wireless LAN (WLAN) networks, and can communicate mission-critical voice, video, and data across peer-to-peer, hierarchical, or meshed networks.

Figure 13-1 Cisco 3200 Series Mobile Access Router

MAR3200 Interfaces

The MAR3200 can be configured with multiple Ethernet and serial interfaces, and up to three radios. The router itself is made up of stackable modules referred to as *cards*. It has two 2.4GHz Wireless Mobile Interface Cards (WMICs) one 4.9GHz WMIC, one Fast Ethernet Switch Mobile Interface Card (FESMIC) and one Mobile Access Router Card (MARC)). Figure 13-2 shows this stackable card configuration. The router can also be configured in a rugged enclosure with power adapters.

Figure 13-2 Card Connections

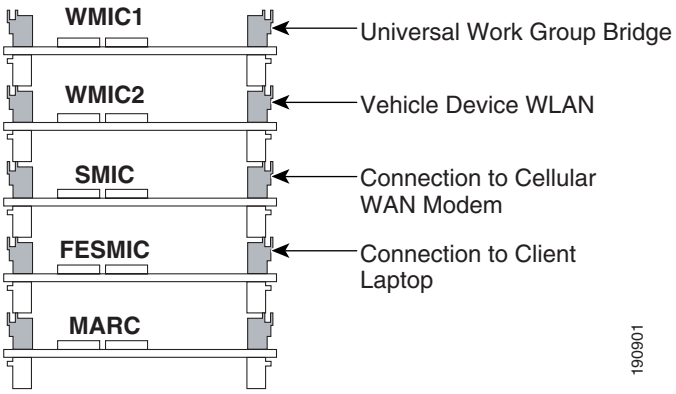

For more information on MAR3200 configuration options, see the following URL:

http://www.cisco.com/en/US/products/hw/routers/ps272/products_data_sheet0900aecd800fe973.html

Figure 13-3 provides an example of a MAR3200 configured with two WMICs, an FESMIC, and a MARC.

Figure 13-3 Mobile Unit Configuration Example

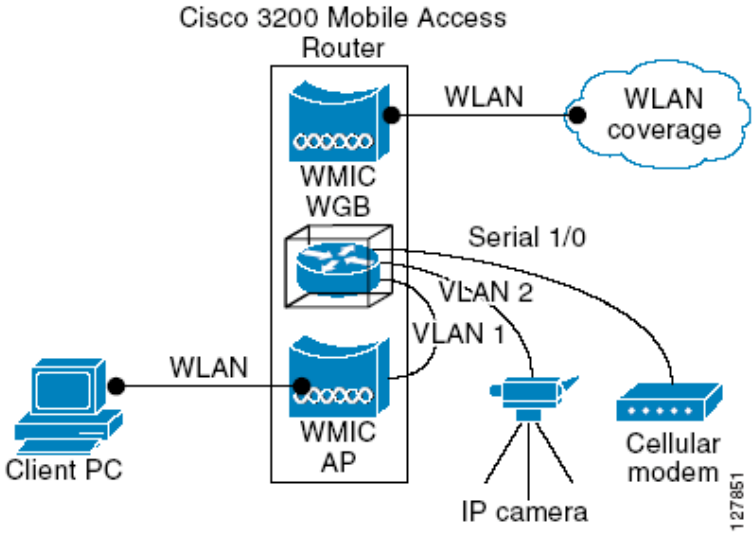

The following tables list the port-to-interface relationships and hardware types. See these tables for configurations where you need to plug other devices into the MAR3200.

Table 13-1 shows the setup of WMICs on the Cisco 3230 Mobile Access router.

Table 13-1 WMIC Ports

	Internal Wiring Ports	Radio Type
WMIC 1 (W1)	FastEthernet 0/0	2.4GHz
WMIC 2 (W2)	FastEthernet 2/3	2.4GHz
WMIC 3 (W3)	FastEthernet 2/2	4.9GHz

Table 13-2 shows the setup of serial interfaces on the Cisco 3230 Mobile Access router.

Table 13-2 SMIC Ports

	Internal Wiring Ports	Interface Type
Serial 0	Serial 1/0	DSCC4 Serial
Serial 1	Serial 1/1	DSCC4 Serial
Internal	Serial 1/2	DSCC4 Serial
Internal	Serial 1/3	DSCC4 Serial

Table 13-3 shows the setup of Fast Ethernet interfaces on the Cisco 3230 Mobile Access router.

Table 13-3 Fast Ethernet Ports

	Internal Wiring Ports	Interface Type
Internal WMIC 1	Fast Ethernet 0/0	Fast Ethernet
FE0X	Fast Ethernet 2/0	Fast Ethernet
FE1X	Fast Ethernet 2/1	Fast Ethernet
Internal WMIC 3	Fast Ethernet 2/2	Fast Ethernet
Internal WMIC 2	Fast Ethernet 2/3	Fast Ethernet

MAR3200 WMIC Features

Table 13-4 highlights the software features of WMICs running Cisco IOS.

Table 13-4 WMIC IOS Software Features

Feature	Description
VLANs	Allows dot1Q VLAN trunking on both wireless and Ethernet interfaces. Up to 32 VLANs can be supported per system.

Table 13-4 WMIC IOS Software Features (continued)

QoS	Use this feature to support quality of service for prioritizing traffic on the wireless interface. The WMIC supports required elements of Wi-Fi Multimedia (WMM) for QoS, which improves the user experience for audio, video, and voice applications over a Wi-Fi wireless connection and is a subset of the IEEE 802.11e QoS specification. WMM supports QoS prioritized media access through the Enhanced Distributed Channel Access (EDCA) method.
Multiple BSSIDs	Supports up to 8 BSSIDs in access point mode.
RADIUS accounting	When running the WMIC in access point (AP) mode you can enable accounting on the WMIC to send accounting data about authenticated wireless client devices to a RADIUS server on your network.
TACACS+ administrator authentication	TACACS+ for server-based, detailed accounting information and flexible administrative control over authentication and authorization processes. It provides secure, centralized validation of administrators attempting to gain access to your WMIC.
Enhanced security	Supports three advanced security features: • WEP keys: Message Integrity Check (MIC) and WEP key hashing CKIP • WPA • WPA2
Enhanced authentication services	Allows non-root bridges or workgroup bridges to authenticate to the network like other wireless client devices. After a network username and password for the non-root bridge or workgroup bridge are set, (LEAP), EAP-TLS or EAP-FAST can be used for authentication in dynamic WEP, WPA, or WPA2 configurations.
802.1x supplicant	In AP mode, the Mobile Access Router supports standard 802.1x EAP types for WLAN clients.
Fast secure roaming	Fast, secure roaming using Cisco Centralized Key Management (CCKM) in Work Group Bridge mode and Universal Work Group Bridge mode.
Universal workgroup bridge	Supports interoperability with non-Cisco APs.
Repeater mode	Allows the access point to act as a wireless repeater to extend the coverage area of the wireless network.

Universal Workgroup Bridge Considerations

The Cisco Compatible eXtensions (CCX) program delivers advanced WLAN system level capabilities and Cisco-specific WLAN innovations to third party Wi-Fi-enabled laptops, WLAN adapter cards, PDAs, WI-FI phones, and application specific devices (ASDs). The 2.4 GHz WMIC provides CCX client support. When the 2.4 GHz WMIC is configured as a universal workgroup bridge client, it does not identify itself as a CCX client. However, it does support CCX features. Table 13-5 lists the supported features.

Table 13-5 CCX Version Feature Support

Feature	v1	v2	v3	v4	AP	WGB	WGB Client
Security							
Wi-Fi Protected Access (WPA)		X	X	X	X	X	X
IEEE 802.11i - WPA2			X	X	X	X	X
WEP	X	X	X	X	X	X	X
IEEE 802.1X	X	X	X	X	X	X	X
LEAP	X	X	X	X	X	X	X
EAP-FAST			X	X	X	X	X
CKIP (encryption)	X				X	X	
Wi-Fi Protected Access (WPA): 802.1X + WPA TKIP		X	X	X	X	X	X
With LEAP		X	X	X	X	X	X
With EAP-FAST			X	X	X	X	X
IEEE 802.11i- WPA2: 802.1X+AE			X	X	X	X	X
With LEAP			X	X	X	X	X
With EAP-FAST			X	X	X	X	X
CCKM EAP-TLS				X	X	X	X
EAP-FAST				X	X	X	X
Mobility							
AP-assisted roaming		X	X	X	X	X	X
Fast re-authentication via CCKM, with LEAP		X	X	X	X	X	X
Fast re-authentication via CCKM with EAP-FAST			X	X	X	X	X
MBSSID				X	X		
Keep-Alive				X	X	X	
QoS and VLANs							
Interoperability with APs that support multiple SSIDs and VLANs	X	X	X	X	X	X	
Wi-Fi Multimedia (WMM)			X	X	X	X	X
Performance and management							
AP-specified maximum transmit power		X	X	X	X	X	X
Recognition of proxy ARP information element (For ASP)			X	X	X		

Table 13-5 CCX Version Feature Support (continued)

Client utility standardization							
Link test				X	X	X	X

MAR3200 Management Options

You can use the WMIC management system through the following interfaces:

- The IOS command-line interface (CLI), which you use through a PC running terminal emulation software or a Telnet/SSH session.
- Simple Network Management Protocol (SNMP)
- Web GUI management

Using the MAR with a Cisco 1500 Mesh AP Network

The Universal Workgroup Bridge feature for the Cisco MAR3200 WMIC allows the WMIC radio to associate to non-Aironet based access points. It also supports a majority of CCXv4 client features. In the version 4.0 software release for the Cisco Wireless LAN Controller (WLC), and Mesh APs, enhancements have been added to support Cisco 1230, 1240, 1130, or 3200 products associating to the Cisco 1500 as a workgroup bridge (WGB). These two feature updates allow the MAR to act as a client to the 1500 Mesh AP networks or Light Weight Access Point Protocol (LWAPP) WLAN networks enabling new solutions for public safety, commercial transportation, and defense markets. The MAR not only has Fast Ethernet and Serial interface connections for other client devices, but can also use them to connect to other network devices for backhaul purposes.

Vehicle Network Example

This section describes a simple application for the MAR3200 in a Mesh network using its universal workgroup bridge feature to connect to the Mesh WLAN. Figure 13-4 illustrates this example.

- A Cisco 3200 Series router installed in a mobile unit allows the client devices in and around the vehicle to stay connected while the vehicle is roaming.
- WMICs in vehicle-mounted Cisco 3200 Series routers are configured as access points to provide connectivity for 802.11b/g and 4.9-GHz wireless clients.
- Ethernet interfaces are used to connect any in-vehicle wired clients, such as a laptop, camera, or telematics devices, to the network.
- Another WMIC is configured as a Universal Workgroup Bridge for connectivity to a Mesh AP, allowing transparent association and authentication through a root device in the architecture as the vehicle moves about.
- Serial interfaces provide connectivity to wireless WAN modems that connect to cellular networks such as CDMA or GPRS. The Wireless 802.11 connections are treated as preferred services because they offer the most bandwidth. However, when a WLAN connection is not available, cellular technology provides a backup link. Connection priority can be set by routing priority, or by the priority for Mobile IP.

Figure 13-4 Vehicle Network Example

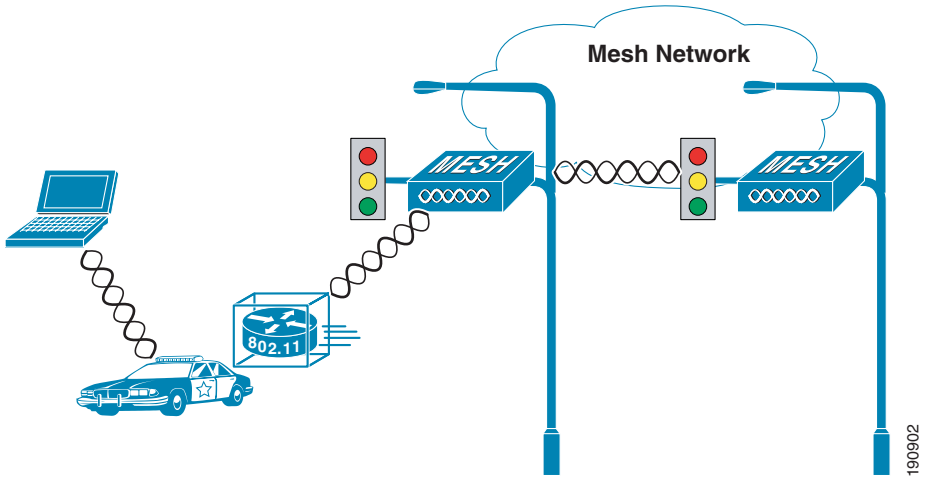

Simple Universal Bridge Client Data Path Example

The IP devices connected to the MAR are not *aware* that they are part of a mobile network. When they must communicate with another node in the network, their traffic is sent to their default gateway, the Cisco 3200 Series router. The Cisco 3200 Series router forwards the traffic to the Mesh APs WLAN, the mesh AP then encapsulates the data packets in LWAPP and forwards them through the network to the controller.

As shown in Figure 13-5, the Cisco 3200 Series router sends traffic over the Universal Bridge Client WLAN backhaul link. This traffic then crosses the WLAN to the controller where it is then forwarded out the controller interface to the wired network. Return traffic destined for any client attached to the MAR would be forwarded via a static route pointing back to the controller of the Mesh network. Figure 13-6 shows the return path to the MAR. Mobile IP eliminates the need for static routing and will be discussed further in this chapter. NAT may be used in simple deployments when Mobile IP is not available.

The data path example shown in Figure 13-5, and previously described, represents the traffic in a pure Layer 2 Mesh when the MAR is using only the WMIC for backhaul. If the deployment calls for more complexity (such as secondary cellular backhaul links) then Mobile IP will be required.

When the WMIC is used as a Universal Bridge Client it sets up its wireless connections the same way any wireless client does.

Figure 13-5 Simple Layer 2 Data Path Example

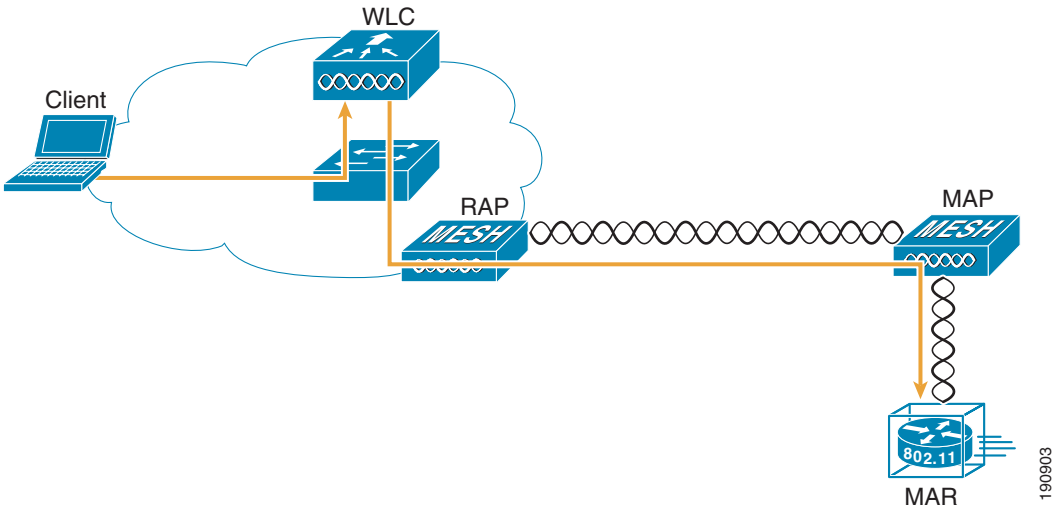

Figure 13-6 Client Return Data Path

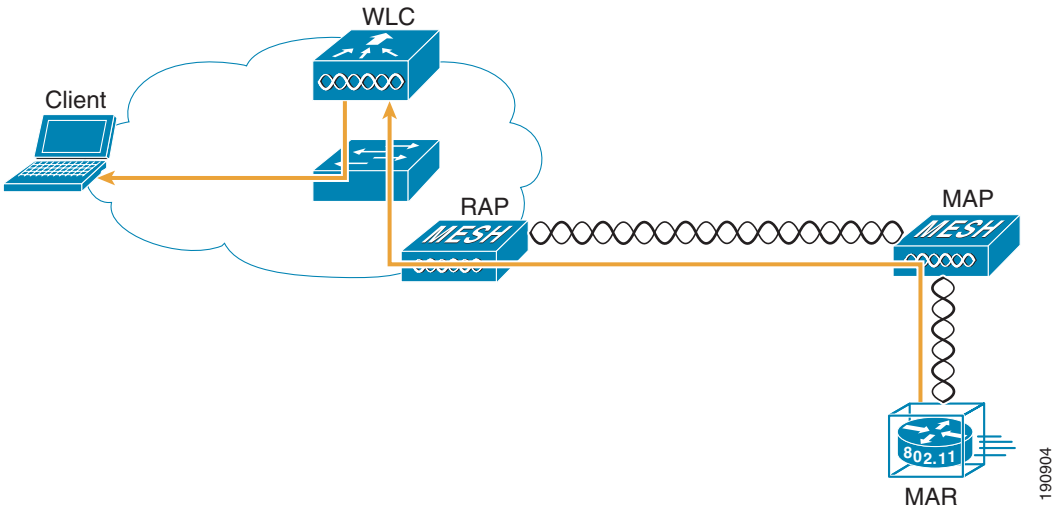

Configuration

The following is a configuration example for the MAR3200. It can be used as a step-by-step process to configure the Universal Work Group Bridge client using open authentication, and WEP encryption. It also covers other basic configuration steps such as VLAN creation, assignment, and DHCP.

Connecting to the Cisco 3200 Series Router

Attach the console cable to both the serial port of your PC and the Mobile Access router console port (DB9 female). Use a straight through DB9-to-DB9 cable.

Note You can also use the same console cable used to access the HA, with the addition of an RJ-45 to DB9 female adapter.

Configuring the IP Address, DHCP, VLAN on MAR

Step 1 Connect to and log into the Mobile Router.

Step 2 Create a loopback interface and assign an IP address.

Step 3 Create VLAN 2 in the VLAN database using the **vlan database** command.

Step 4 Configure the VLAN 3 and VLAN 2 interfaces. VLAN 3 is used for the 2.4 GHz WMIC2 (W2) which is acting as AP and VLAN 2 is used for the 4.9GHz WMIC (W3). Configure FA2/0, FA2/1 and FA2/3 to be in VLAN 3, and FA 2/2 to be in VLAN 2.

Step 5 Create VLAN 4 in the VLAN database for connection between WMIC 1 and MARC.

Table 13-6

Connected to	Interface	Radio Type	VLAN	Description
PC	FastEthernet2/0	None	3	Fast Ethernet link for end device
WMIC 1 (W1)	FastEthernet0/0	2.4GHz	4	2.4 GHz Universal Work Group Bridge connection to Mesh Network
WMIC 2 (W2)	FastEthernet2/3	2.4GHz	3	Provide 2.4 GHz AP Hotspot around mobile router
WMIC 3 (W3)	FastEthernet2/2	4.9GHz	2	4.9GHz uplink as Workgroup Bridge

Step 6 Configure DHCP server for VLAN 3 using following commands:

ip dhcp pool mypool

network 10.40.10.0 /28

default-router 10.40.10.1

ip dhcp excluded-address 10.40.10.1 10.40.10.3

Step 7 Verify that the wired client on VLAN 3 has been assigned a DHCP IP address in the **10.40.10.0/28** subnet.

Configuring the Universal Bridge Client on WMIC

This configuration is made on the WMIC, and is used for connecting the Mobile Access router (MAR) to a Cisco Mesh network.

Step 1 Configure the SSID of the mesh network on the MARs WMIC with which you plan to connect.

Step 2 Connect to the console port of the WMIC:

dot11 ssid (A given SSID)

Step 3 Configure your authentication type:

authentication (Auth Type)

client EAP client information

key-management key management

network-eap leap method

open open method

shared shared method

Step 4 Configure your encryption key, if needed:

encryption key 1 size 128bit 7 FA1E467E23EAD518A21653687A42 transmit-key

encryption mode wep mandatory

Step 5 Configure the WMIC to act as a universal client to the Mesh network:

station-role workgroup-bridge universal (mac address)

Note You must use the MAC address of the associated VLAN that the WMIC is bridged to. For example, to use the MAC address of VLAN 1, acquire the MAC address of VLAN 1 by entering the **show mac-address-table** command from the console of the MARs router card.)

Step 6 Bridge the dot11 interface:

bridge-group 1

bridge-group 1 spanning-disabled

Step 7 Bridge the ethernet interface:

FastEthernet0

bridge-group 1

Step 8 Configure the bridged virtual interface:

interface BVI1

no ip address

no ip route-cache

Configuring the MARs Router Card

The following configuration is for the router card of the MAR.

Step 1 Find the interface the WMIC is associated with by issuing the following command:

show CDP neighbors!

Step 2 Configure the interface with the matching VLAN that you used in Step 5 for the MAC address in the universal client command:

interface FastEthernet2/2

switchport access vlan 4

Step 3 Configure the VLAN to use DHCP if you are going to be using DHCP on the MAR:

interface Vlan4

ip address dhcp

WMIC Roaming Algorithm

Four basic triggers start the WMIC scanning for a better root bridge or access point:

- The loss of eight consecutive beacons
- The data rate shifts
- The maximum data retry count is exceeded (the default value is 64 on the WMIC)
- A measured period of time of a drop in the signal strength threshold

Only the last two items in this list are configurable using the **packet retries** command and **mobile station period** *X* **threshold** *Y* (in dBm); the remainder are hard-coded.

If a client starts scanning because of a loss of eight consecutive beacons, the message "Too many missed beacons" is displayed on the console. The WMIC in this case acting as a universal bridge client much like any other wireless client in its behavior.

An additional triggering mechanism, mobile station, is not periodic but does have two variables: *period* and *threshold*.

If mobile station is configured. The mobile station algorithm evaluates two variables: data rate shift and signal strength and responds as follows:

- If the driver does a long term down shift in the transmit rate for packets to the parent, the WMIC initiates a scan for a new parent (no more than once every configured period).
- If the signal strength (threshold) drops below a configurable level, the WMIC scans for a new parent (no more than once every configured period).

The data-rate shift can be displayed with the **debug dot11 dot11Radio 0 trace print rates** command. However, this will not show the actual data rate shift algorithm in action, only the changes in data rate. This determines the time period to scan depending on how much the data rate was decreased.

The period should be set depending on the application. Default is 20 seconds. This delay period prevents the WMIC from constantly scanning for a better parent if, for example, the threshold is below the configured value.

The threshold sets the level at which the algorithm is triggered to scan for a better parent. This threshold should be set to *noise+20dBm* but not more than -70dBm (+70 since input for threshold is positive). The default is -70 dBm.

MAR3200 in a Mobile IP Environments

The wireless technologies used in many current metropolitan mobile networks include 802.11 wireless mesh networks for general city-wide coverage, providing high speed access for bandwidth-intensive applications, such as in-car video. For coverage areas where it is not practical to extend the wireless mesh network, it can be supplemented by cellular services, such as CDMA 1x RTT. Using this approach, cellular services can be used to fill gaps in connections and provide backup wireless connectivity. This added backup interface requires Mobile IP to enable client roaming between the two separate networks.

To enable Mobile IP, a Home Agent (HA) router must be added to the enterprise network to tunnel client traffic between the Mobile Router and its home network. Another requirement for Mobile IP is to configure the MAR3200 as a Mobile Router (MR). The following section describes Mobile IP registration process. Figure 13-5 displays a very simple Mobile IP (MIP) environment.

MAR 3200 Mobile IP Registration Process

When the MAR3200 is associated to its Mesh network, the following events occur:

- The MAR3200 goes through a Foreign Agent (FA) discovery process.

  FAs advertise their existence periodically. If a MR does not hear a FA advertisement, it solicits itself by sending a multicast advertisement to the address 224.0.0.2.

- If an FA receives a solicitation from an MR, it responds with a unicast advertisement to the MR that includes its Care of Address (CoA).
- If the access network does not have a FA router, the MR can register itself with the HA by using a Collocated Care of Address (CCoA).

  The CCoA address is the IP address of the interface the MR uses to connect to the access network.

- The MR then sends in Registration Request (RRQ) to the HA.
- The HA authenticates MR by sending a Registration reply (RRP) to the MR.
- The HA provides a gratuitous APR update for the home network, then creates a GRE tunnel to the FA if using Foreign Agent CoA (FACoA), or to the MR if you are using CCoA. It then adds a host route to the MR.
- Now, the MR has reached a registered state with the HA and the HA has set up a binding table entry for the MR CoA. It will then tunnel and route traffic destined for the MR.
- At this point, the mobile router is registered through a Mesh WLAN to its HA using the FACoA.

  If any devices attached to the Cisco 3200 Series router must communicate with nodes on the home network, they send the data to the Cisco 3200 Series router and Mobile IP tunnels the data to the HA, with any traffic directed to MR clients tunneled from the HA to the MR. A simple Mobile IP network with FACoA for Mesh and Collocated Care of Address (CCoA) for cellular is illustrated in Figure 13-7. Mobile IP is needed if your application requires routing to any devices or nodes attached to the MAR3200.

- If the MAR3200 is not in the vicinity of a wireless LAN hot spot it can use a backup wireless service such as cellular modem to deliver the data.

  In this case, the Cisco 3200 generates a CCoA from the IP address it acquired from the service provider network and registers its CCoA with the home agent. This CCoA address is the mobile router's own interface IP address it acquired via DHCP from the Service Provider. The registration process is similar to the process for CoA registration.

Figure 13-7 Mobile IP Example

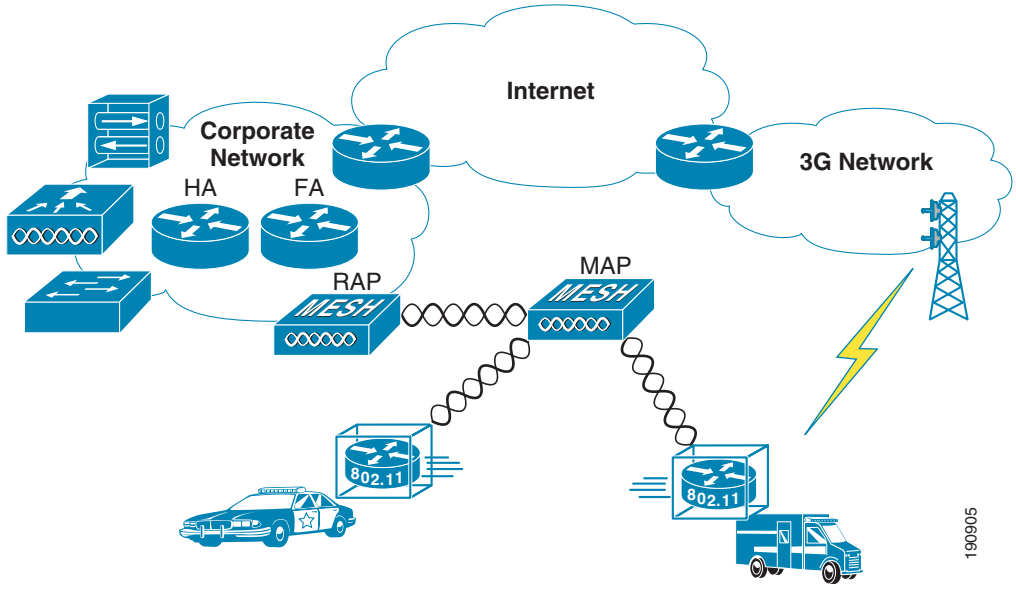

For more information on Mobile IP, see the following URL:
http://www.cisco.com/en/US/tech/tk827/tk369/tk425/tsd_technology_support_sub-protocol_home.html

CHAPTER 14

Cisco Unified Wireless and Mobile IP

Introduction

This chapter describes the inter-workings of the Cisco Mobile Client (CMC) over a Cisco Unified Wireless Network (WiSM). This chapter covers the following topics:

- Different levels of mobility
- Requirements for a mobility solution
- Roaming on the Cisco Unified Wireless Network
- Roaming on a Mobile IP-enabled network
- CMC characteristics when roaming on a Cisco Unified Wireless Network

Different Levels of Mobility

There are three different levels of mobility:

- Layer 2 roaming across a single Layer 2 network:
  - All of the APs are on the same subnet without trunking
- Layer 3 roaming across a single Layer 2 network:
  - Cisco Unified Wireless Network
- Layer 3 roaming across any Layer 2 network:
  - Mobile IP Client

One example of Layer 2 roaming across a single Layer 2 network is a wireless network where all the APs have to be on the same subnet and the clients roam between them. This type of deployment allows the clients to roam from one AP to a new AP without requiring a new IP address or the network being mobility-aware.

Layer 3 roaming across a single Layer 2 network follows the previous AP example, but allows the APs to be on different subnets while also allowing the clients to remain in the same subnet as they roam from AP to AP. Layer 3 roaming across any Layer 2 network is a generalized version of this concept to allow roaming across completely different Layer 2 networks (cellular, wired, and 802.11 wireless).

Roaming in networks can be seamless roaming or seamless *mobility.* Seamless mobility is where both the mobile client applications and the remote applications do not notice any change in end-to-end IP addressing; end applications can use or embed these IP addresses into their data packets without concern that they will be undeliverable. This emulates the case where two clients are on a wired network and not mobile. The Cisco Unified Wireless Network and Mobile IP both provide seamless mobility.

The Cisco Unified Wireless Network is an example of seamless Layer 3 roaming across a single Layer 2 network, while the CMC using Mobile IP (RFC 3344) is an example of seamless Layer 3 roaming across any Layer 2 network. That is, in the Cisco Wireless Unified Network, Layer 3 roaming is restricted to roaming across APs in the mobility group. With Mobile IP, any Layer 2 network (wired, 802.11 wireless, or cellular) can be used for roaming.

These two different solutions perform the same functionality, so they require the same components.

Requirements for a Mobility Solution

There are the following five requirements for every mobility solution:

- Location database
- Move discovery
- Location discovery
- Update signaling
- Path re-establishment

These requirements are covered in the following sections.

Location Database

A location database keeps track of the roaming client. This is very important because the location database is actually forwarding all packets destined for the client to the current location of the roaming client. That is, the location database receives packets destined for the client and forwards the packets on to the client.

In the Cisco Unified Wireless Network, the first hop router attracts packets for the wireless clients through the routing protocol running on that network, and forwards them via a trunk to the controller. Each controller keeps a location database of wireless clients as they roam from one AP to another AP associated to the controller. If the wireless client then roams to an AP on another controller (a foreign controller), that controller can query other controllers in the mobility group to see if this is a new client or a roaming client. If it is a roaming client, the first hop router near the home controller still attracts packets destined to the wireless client, but instead of the controller forwarding them on to one of its associated APs, it forwards the packets to the foreign controller, which then forwards them on to the client.

In Mobile IP, the Home Agent (HA) is the location database. Because it runs the network routing protocol, it attracts packets for the Mobile IP Client and forwards them to the current location of the client. Unlike the Cisco Unified Wireless Network, the HA is not a distributed database between WLCs. It does not query other HAs. As far as it is concerned, there is only one location database: itself. This is where the location database mechanisms for the two solutions differ.

Move Discovery, Location Discovery, and Update Signaling

The discovery, location discovery, and update signaling requirements are grouped in this section because in the Cisco Unified Wireless Network, they are performed at the same time. When the wireless client roams to a new AP, it needs to associate to the wireless network. During the association process, packets are sent to the controller to identify the wireless client and the location (AP) from where the wireless client is trying to associate. This information is used by the controller to update its mobility database. If the client has roamed to another controller, the original controller for the wireless client forwards packets destined to the wireless client to the remote controller.

Move discovery is done in the Cisco Unified Wireless Network by the network that knows to which AP the wireless client is associated. Update signaling is done by the first packets sent to the controller from the wireless client. These packets can be authentication packets.

In Mobile IP, the Mobile IP Client authenticating to the wireless network does not provide the HA with any information. Additionally, the client is responsible for recognizing when it has moved. The client typically detects movement in two ways. One way is through the Windows operating system's Layer 2 notification feature called Media Sense. This feature detects the disconnect and reconnect of different Layer 2 media when roaming between APs and sends the Windows operating system a signal when it occurs. This allows the interface to try and renegotiate its DHCP address with the DHCP server.

The second method for detecting movement is through FA advertisements. These advertisements tell the Mobile IP Client which subnet it is on. If the Mobile IP Client receives one of these periodic messages, it can tell it has moved to a new subnet. These move discovery methods are typically used for Mobile IP. There are other methods specified in RFC 3344, but generally these are not used in working clients. The next section explains location discovery for Mobile IP.

Location discovery is typically done in one of two ways in Mobile IP. In the first method, it receives an FA advertisement telling it what the IP address is for the FA. The Mobile IP Client can check this address against the address it already has from the FA and tell if the FA has changed locations. The Mobile IP Client can then forward this IP address to its HA so that the HA can forward packets to the Mobile IP Client. In the second method, the client is acting as its own FA, it receives a new DHCP IP address, and informs the HA it has a new address for forwarding packets.

Finally, update signaling in Mobile IP is done via the registration request (RRQ) and registration reply (RRP) between the Mobile IP Client and the HA. These packets have a cryptographic signature (via shared keys) to ensure that they are not changed in transit.

For more information, see the following URL:

http://www.cisco.com/go/mobile_ip

Path Re-establishment

Path re-establishment is the mechanism used to allow the client to receive packets that are destined for it from the location database. This is typically some type of tunneling where the original packet is encapsulated into another packet.

In the Cisco Unified Wireless Network, packets are forwarded to wireless clients on associated APs through the "always up" LWAPP tunnel. For wireless clients that have roamed to another controller, the controllers use a dynamic Ethernet over IP tunnel for all packets forwarded to other controllers in the mobility group.

In Mobile IP, there are several types of tunnels available (GRE, UDP, and IP in IP) and the type of tunnel used depends on the equipment between the Mobile IP Client and HA, and whether the HA supports that type of encapsulation. For example, if the HA detects that the client is behind a NAT gateway, it uses

UDP tunneling. If the Mobile IP Client requests GRE tunneling and the HA can support the tunneling, it uses GRE. Typically, the Mobile IP Client requests IP in IP tunneling, and all RFC-compliant HAs can support this type of tunneling.

Roaming on a Cisco Unified Wireless Network

A Cisco Unified Wireless Network acts as a mobility proxy for the wireless client. This allows the network to provide seamless mobility to the wireless client without any extra software or required information on the wireless client (see Figure 14-1).

When a wireless client associates to an AP, the AP forwards the client packets to the controller via the LWAPP tunnel set up between the controller and AP (the LWAPP tunnel is set up between the AP and controller at AP boot time). For the controller, the LWAPP tunnel allows it to do the following:

- Know to which AP the client is associated (LWAPP tunnel endpoint)
- Forward packets back to the client via the tunnel
- Be multiple hops away from the AP and still receive the client traffic
- Filter the packets to and from the wireless client

For the client, the LWAPP tunnel allows the client to see its default gateway as being one hop away, even though it might physically be several hops away.

Figure 14-1 Roaming on a Cisco Unified Wireless Network

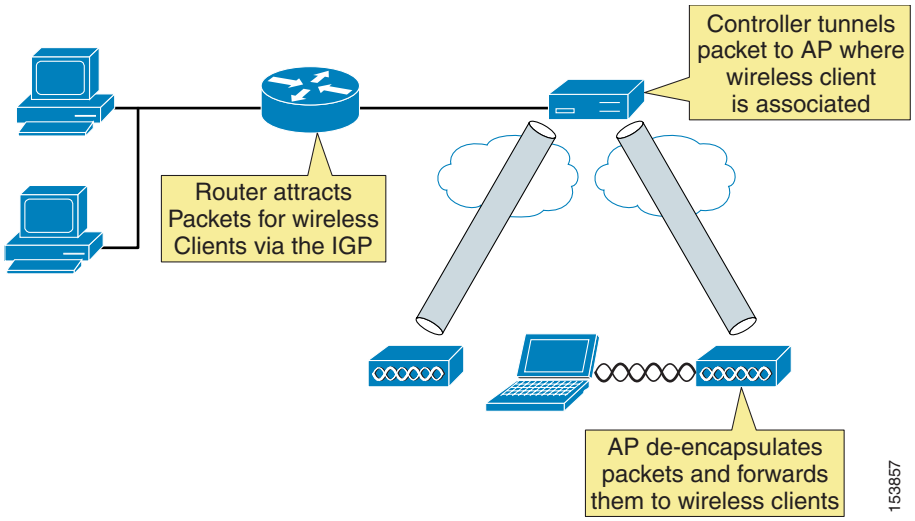

If the client requests a DHCP address, the controller either gives the client an address from its local DHCP pool (if defined) or fills in the gateway address in the DHCP request for an external DHCP server. In either case, the controller modifies any returning offers so that the DHCP server's address is set to the address on its virtual interface. Even though the virtual IP address is not in any routing table (typically 1.1.1.1), it still allows the controller to intercept any DHCP renewals on wireless clients that occur with the Microsoft Windows operating system (using Microsoft Media Sense) when it roams between APs. In addition, if the same address is on all controllers' virtual addresses, it allows other controllers to intercept the DHCP renewal from the client when it roams to a new AP associated to a different controller.

The wireless client can easily roam between any APs associated to the controller because the controller simply keeps track of the wireless client's current location and forwards the packets destined to that client into the correct LWAPP tunnel and on to the associated AP. When the client roams to an AP associated to a different controller, the remote controller queries other controllers in the mobility group to see if the client has roamed from another controller, and the controllers dynamically set up an Ethernet over IP tunnel for forwarding client traffic from the original controller.

Traffic originating from the wireless client that has roamed to an AP associated to another controller can be handled in two ways. Typically, the foreign controller modifies the destination MAC address of any packet from the wireless client to its gateway MAC address before forwarding it on to the controller gateway. The second method occurs if mobility anchoring is enabled on the original controller; the traffic is forwarded back to the original controller. This allows traffic to be sent to the correct gateway in case RPF checks are enabled.

Roaming on a Mobile IP-enabled Network

A Mobile IP-enabled network has three components:

- Mobile node (MN)—Mobile IP Clients (notebooks)
- Home Agent (HA)—Serves as location database for MNs and attracts the MNs packets by advertising reachability to the MN in the Interior Gateway Protocol (IGP). Tunnels packets to MN.
- Foreign Agent (FA)—(Optional) Offloads CPU processing of encapsulation and decapsulation from the MN and saves IP address space. FAs are not often deployed in enterprise campus environments.

Only two of the three components are actually required for a mobility solution: the MN and HA. The third component, the FA, is optional because the MN can act as its own FA by using DHCP for a local IP address. In this case, the tunnel ends at the MN.

In Figure 14-2, the MN is given an IP address (10.1.69.10) local to the HA. To the rest of the network, the MN looks like it is directly attached to the HA. The HA then uses its mobility binding table to forward packets to wherever the MN is currently located. It is the responsibility of the MN to update its location with the HA. The FA de-encapsulates the packets destined for the MN and forwards them out its interface. It gleans the information it needs by being an active party in the registration process with the HA. The MN actually sends its packets to the FA, and the FA checks the packets and generates new IP headers to forward the information onward to the HA. The FA can also provide reverse tunneling for the MN originated packets back to the HA, instead of simply forwarding through the normal switching process. Reverse tunneling allows packets from the MN to always exit the HA and pass any reverse path forwarding (RPF) checks.

Figure 14-2 HA and FA Tunneling

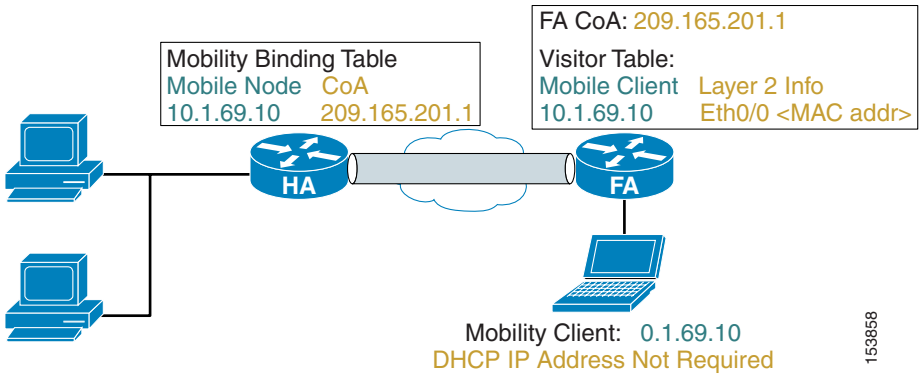

Unlike the Cisco Unified Wireless Network where the network proxies or provides the wireless client with seamless mobility and no information is stored on the client, the Mobile IP Client, or MN needs to know three pieces of information to function:

- Its home address (on a locally connected subnet on the HA)
- Its HA address (so it can update the HA with its current location)
- Its shared secret key (used to authenticate packets between the MN and HA)

These three pieces of information can be dynamically discovered or generated but are typically manually configured on the MN. DHCP can be used to convey the HA address to the MN via option 68. The HA can dynamic assign an IP address to the MN to be used as its home address when it registers for the first time. By using the Cisco Zero Configuration Client (ZECC) feature, the MN and HA can automatically generate a shared secret key from the windows login credentials.

When the Cisco Mobile IP Client is loaded on a Windows host, the Mobile IP Client function rests between the physical interfaces and the VPN client and TCP/IP stack (see Figure 14-3). The Mobile IP Client function sends its home address up the TCP/IP stack so that the host applications, including the VPN client, see a constant IP address as the MN roams across the different network locations or different networks. The physical interfaces might or might not have IP addresses during roaming depending on whether an FA is present on the subnet.

Figure 14-3 Mobile IP Function Position in the Microsoft Operating System

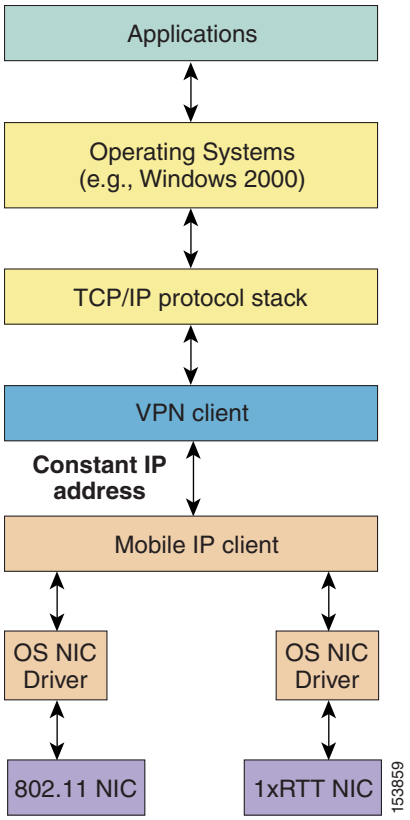

The CMC controls that interface with host-originated packets are transmitted by:

- Installing a new virtual interface adapter (CMIPDRV) at install time.
- Modifying the host forwarding table.

This virtual adapter looks like any physical adapter to the host (see the example in Sample Mobile IP Client Interface and Host Table Manipulation, page 14-8). When the adapter is enabled, the Mobile IP Client modifies the forwarding table to give the CMIPDRV adapter the best metric, and the Windows operating system forwards host originated packets to the CMIPDRV adapter. This allows the Mobile IP Client to hide the true interface used to transmit the packet and to modify the host's forwarding behavior.

In the example, there are three interfaces:

- A local area connection with a static IP address and no gateway.
- A Mobile IP Client interface (CMIPDRV) with a configured home address and gateway
- A wireless connection that has an address filled in by Mobile IP as 0.0.0.0. The actual address is not shown to the Windows operating system.

Not e that the Mobile IP Client has manipulated the host's forwarding table so that the lower metric interface is the Mobile IP Client's interface. The higher metric routes can be safely ignored when looking at the table. The real DHCP IP address on the wireless interface is 10.20.41.12. Any route with a destination address to this gateway has had its metric raised and the default gateway is via the CMIPDRV interface.

Sample Mobile IP Client Interface and Host Table Manipulation

```
C:\>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:

  Connection-specific DNS Suffix . :
  IP Address . . . . . . . . . . . . : 10.20.30.249
  Subnet Mask. . . . . . . . . . . . : 255.255.255.0
  Default Gateway  . . . . . . . . . :

Ethernet adapter CMIPDRV:

  Connection-specific DNS Suffix . : srnd3.com
  IP Address . . . . . . . . . . . . : 10.20.32.11
  Subnet Mask. . . . . . . . . . . . : 255.255.255.0
  Default Gateway  . . . . . . . . . : 10.20.32.1

Ethernet adapter Wireless Connection:

  Connection-specific DNS Suffix . :
  IP Address . . . . . . . . . . . . : 0.0.0.0
  Subnet Mask. . . . . . . . . . . . : 0.0.0.0
  Default Gateway  . . . . . . . . . :

C:\>route print
===========================================================================================
Interface List
0x1...........................MS TCP Loopback interface
0x2...00 d0 b7 a6 b8 47........Intel (R) 82559 Fast Ethernet LAN on Motherboard
- Packet Scheduler Miniport
0x3...00 4d 69 70 56 61 .......Cisco Systems Mobile Adapter - Packer Scheduler
Miniport
0x10005...00 12 f0 7c a5 ca......Intel (R) PRO/Wireless 2915ABG Network Connec
tion - Deterministic Network Enhancer Miniport
===========================================================================================
===========================================================================================
Active Routes:
Network Destination           Netmask         Gateway       Interface      Metric
            0.0.0.0           0.0.0.0      10.20.32.1     10.20.32.11           1
         10.20.30.0     255.255.255.0    10.20.30.249    10.20.30.249           1
         10.20.30.0     255.255.255.0      10.20.32.1     10.20.32.11           1
       10.20.30.249   255.255.255.255       127.0.0.1       127.0.0.1           1
         10.20.32.0     255.255.255.0     10.20.32.11     10.20.32.11          20
        10.20.32.11   255.255.255.255       127.0.0.1       127.0.0.1          20
         10.20.41.0     255.255.255.0     10.20.41.12     10.20.41.12          25
         10.20.41.0     255.255.255.0      10.20.32.1     10.20.32.11           1
        10.20.41.12   255.255.255.255       127.0.0.1       127.0.0.1          25
     10.255.255.255   255.255.255.255    10.20.30.249    10.20.30.249           1
     10.255.255.255   255.255.255.255     10.20.32.11     10.20.32.11          20
     10.255.255.255   255.255.255.255     10.20.41.12     10.20.41.12          25
          127.0.0.0         255.0.0.0       127.0.0.1       127.0.0.1           1
          224.0.0.0         240.0.0.0    10.20.30.249    10.20.30.249           1
          224.0.0.0         240.0.0.0     10.20.32.11     10.20.32.11          20
          224.0.0.0         240.0.0.0     10.20.41.12     10.20.41.12          25
    255.255.255.255   255.255.255.255    10.20.30.249    10.20.30.249           1
    255.255.255.255   255.255.255.255     10.20.32.11     10.20.32.11           1
    255.255.255.255   255.255.255.255     10.20.41.12     10.20.41.12           1
Default Gateway:           10.20.32.1
===========================================================================================
Persistent Routes:
  None
```

When an MN makes a Layer 2 connection, it starts two different threads. One thread is a DHCP process to obtain a local IP address so that it can use the IP address for a co-located care of address (CCoA) registration to the HA if there is no Foreign Agent (FA) on the subnet; the other thread looks for a FA on the subnet to which it is attached.

If the MN finds an FA on the subnet, it uses the care of address (CoA) advertised by the FA to register (update) with its location database, the HA, and reject any DHCP offers. An FA on the subnet does two things for the Mobile IP Client:

- The HA forms a tunnel with the FA CoA to forward packets destined for the MN, thereby relieving the MN of having to obtain a local address. The FA forwards packets to the MN home address on its local interfaces via Layer 2 information it gleaned during registration with the HA.
- It offloads the tunnel packet processing of encapsulation or de-encapsulation to the FA.

The FA can forward traffic to the MN because the MN is on a directly attached interface. The FA maintains an entry in a table, called a visitor table, which has the MN home address, and to which interface the MN is currently attached as well as Layer 2 encapsulation information. This way, when the HA tunnels a packet for the MN to the FA, the FA simply de-encapsulates the packet and looks into its visitor table for the interface the MN is on and forwards it directly out the interface. Because of this table, the MN does not need a local IP address on the subnet.

If there is no FA on the subnet, the MN requires a local IP address to which the HA can forward packets. After it receives a DHCP address, the MN registers (updates) the HA and builds a tunnel directly between the MN and the HA. All de-encapsulation of packets is performed by the MN.

If reverse tunneling (where the host packets are tunneled back to the HA) is enabled, the overall solution is analogous to the Cisco Unified Wireless Network. Packets from the client are tunneled and forwarded to a location database and packets destined to the client are received by the location database and tunneled and forwarded to the current location of the client.

Figure 14-2 and Figure 14-3 are similar in functionality except that the HA is a router and can also attract packets for the Mobile IP Client through the use of an IGP and tunnel packets to the MN.

Cisco Mobile IP Client Characteristics When Roaming on a Cisco Unified Wireless Network

Traffic destined for the MN must pass through the HA and the controller to reach a MN on the wireless network. If reverse tunnel is enabled, the packet must pass back through the HA before being forwarded to any other host. Figure 14-4 shows the traffic patterns from a remote host to the MN. The orange flow line shows that the network believes the MN is attached to the HA. The blue flow line shows the tunneled packet to the MN. If another wireless client sent packets to the MN, that traffic would also have to traverse the HA.

Figure 14-4 Traffic Flow to MN

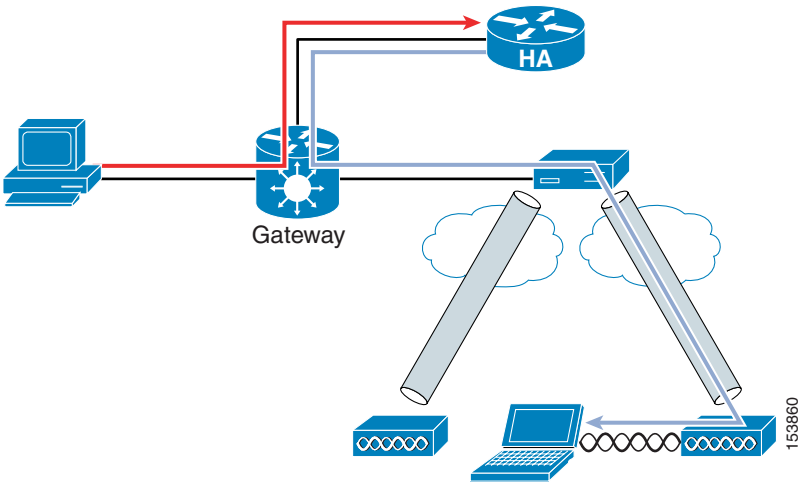

Because of the routing of traffic to and from the MN from other hosts, the general goal in the placement of the controller and HA is to minimize the summation of all links. In Figure 14-5, link 1 cannot be minimized because the hosts' locations are random. The same goes for link 4 because mobile hosts' locations cannot be fixed. Link 3 cannot be minimized because the RF survey determines AP placement. This leaves the link between the controller and HA, link 2.

Figure 14-5 Mobile IP and Cisco Unified Wireless Network

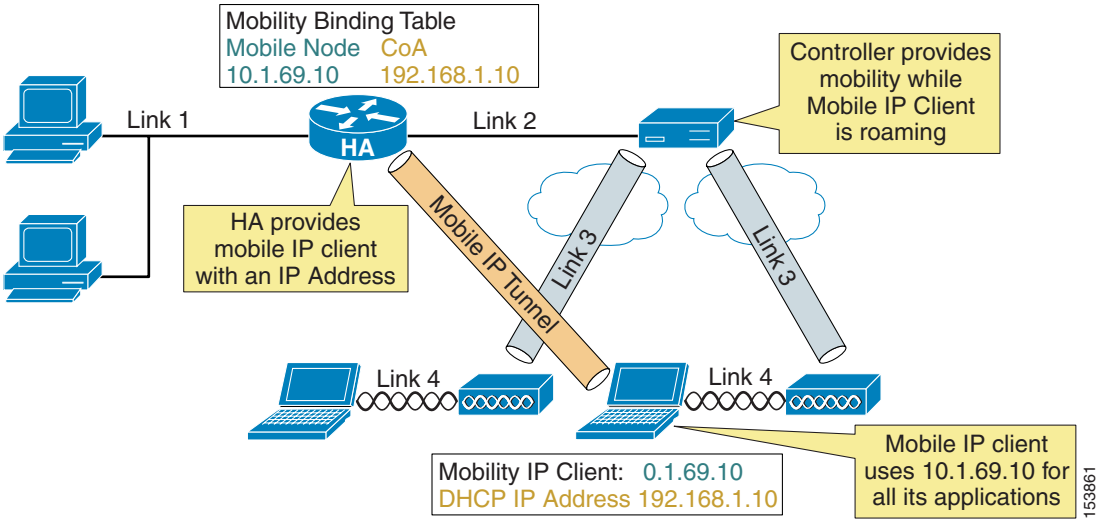

There are two basic HA placement principles:

- HA placement must be as close to the core as possible
- HA placement must be as close to the controller as possible

The first principle is simply a way to minimize traffic links from any host in the network to any place in the network. The second principle follows the logic that the only link you can minimize is link 2 between the HA and controller. This means the controller and HA should be co-located whenever possible. The best location is directly off the core in the data center with centralized controllers.

When a Mobile IP Client is roaming on a Cisco Unified Wireless Network, it maintains the same DHCP IP address while roaming, allowing it to maintain the same CCoA address. The Cisco Unified Wireless Network handles the underlying mobility and the Mobile IP Client does not see any changes as it roams from AP to AP. To the Mobile IP Client, it is as if it is roaming on a single large subnet. Accordingly, nothing changes at the Mobile IP Client level until it roams off of the wireless network.

Note CCoA mode for the CMC is recommended on the Cisco Unified Wireless Network because of unwanted multicast traffic over the shared wireless network when multicast is enabled at the controller. Because multicast traffic is disabled at the controller by default, there is no requirement for FAs on the wireless network. See Chapter 6, "Cisco Unified Wireless Multicast Design," for more information about the multicast traffic on a Cisco Unified Wireless Network.

Currently, the CMC does not behave in this manner. When it roams from AP to AP, even though it retains the same DHCP address as it roams, with each roam, it reregisters with the HA, as though it received a new DHCP IP address. Although this is a minor nuisance, it still taxes the resources of the home agent.

Another characteristic of the Mobile IP Client is that it runs in parallel with the Microsoft Windows operating system. When the Mobile IP Client controls the DHCP behavior of the interface instead of the Windows TCP/IP stack, the Windows operating system gives a connectivity warning for the interface indicating that the interface is configured for DHCP but did not receive an IP address. This is normal operation in Mobile IP but the Windows operating system considers it a warning situation. In addition, the Mobile IP Client configures the interface with a 0.0.0.0 IP address. This can be confusing to users troubleshooting the connectivity problem brought up by the Windows operating system.

Another characteristic of the CMC is that it does not yet use the Cisco Zero Configuration Client (ZECC) in version 1.0. For more information on ZECC, see the following URL:

http://www.cisco.com/en/US/partner/tech/tk827/tk369/technologies_white_paper0900aecd8021a77d.shtml)

This forces network engineers to configure users in two places:

- On a domain server to allow clients to log into the domain
- On a AAA server where the MN's shared secret keys are stored

Version 1.0 of the CMC requires a RADIUS server for storage of MN keys (keys can also be locally stored on the HA but this does not scale). Cisco recommends Cisco Secure Access Control Server (ACS) 4.0. For information on how to store keys on the RADIUS server, see the following URL:

http://www.cisco.com/en/US/partner/tech/tk827/tk369/technologies_white_paper09186a00801fe71d.shtml

One final note: The current version of the roaming server does not have the ability to configure the Cisco HA, and the HA is considered an external HA. This limits the roaming server to a RADIUS database in the CMC architecture. The roaming server is a GUI-based AAA server that can also configure ipUnplugged's HA. Future versions of Roaming server will allow configuration of Cisco HA.

CHAPTER 15

Cisco Unified Wireless Location-Based Services

Introduction

With integrated location tracking, enterprise wireless LANs become much more valuable as a corporate business asset. Enterprise network administrators, security personnel, and others directly responsible for the health and well-being of business-class networks have expressed great interest in location-based services to allow them to better address issues in their environments such as the following:

- Quickly and efficiently locating valuable assets and key personnel
- Improving productivity via effective asset and personnel allocation
- Reducing loss because of the unauthorized removal of assets from company premises
- Improving customer satisfaction by rapidly locating critical service-impacting assets.
- Improving WLAN planning and tuning capabilities.
- Coordinating Wi-Fi device location with security policy enforcement
- Meeting regulatory requirements for E911 calls

This chapter discusses the Cisco Location-Based Service (LBS) solution and the areas that merit special consideration involving design, configuration, installation, and deployment. Each of these areas is described in brief and reference is made to a comprehensive white paper entitled *Wi-Fi Location-Based Services: Design and Deployment Considerations*, which contains in-depth discussion and analysis and is available at the following URL: http://www.cisco.com/univercd/cc/td/doc/solution/wifidesi.pdf. This chapter addresses the following topics:

- The fundamentals of positioning technologies including lateration, angulation, and location patterning approaches.
- Cisco RF Fingerprinting and its advantages over traditional positioning techniques
- Traffic flow analysis between the Cisco Wireless Location Appliance and other network components
- In-depth discussion of various RFID tag technologies including vendor-specific configuration information
- How external third-party location client applications can interface with the Cisco Wireless Location Appliance

Reference Publications

This document makes extensive reference to the following white paper, which should be referenced for further detailed information regarding any section in this chapter:

- Wi-Fi Location-Based Services: Design and Deployment Considerations—http://www.cisco.com/univercd/cc/td/doc/solution/wifidesi.pdf

Additionally, review the following supplemental documents:

- Release Notes for Cisco Wireless Location Appliance—http://www.cisco.com/en/US/products/ps6386/prod_release_note09186a00806b5ec7.html
- Cisco Wireless Location Appliance: Installation Guide—http://www.cisco.com/en/US/products/ps6386/products_installation_and_configuration_guide_book09186a00804fa761.html
- Cisco Wireless Location Appliance: Configuration Guide—http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_book09186a00806b5745.html
- Cisco Wireless Location Appliance: Deployment Guide—http://www.cisco.com/en/US/products/ps6386/prod_technical_reference09186a008059ce31.html
- Cisco Wireless Control System Release Notes, Release 4.0—http://www.cisco.com/en/US/products/ps6305/prod_release_note09186a00806b0811.html
- Cisco Wireless Control System Configuration Guide, Release 4.0—http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_book09186a00806b57ec.html

In addition, Cisco recommends that you review Chapter 8, “Cisco Unified Wireless Control System,” in this design guide.

Cisco Location-Based Services Architecture

Positioning Technologies

Location tracking and positioning systems can be classified by the measurement techniques they employ to determine mobile device location (*localization*). These approaches differ in terms of the specific technique used to sense and measure the position of the mobile device in the target environment under observation. Typically, *Real Time Location Systems (RTLS)* can be grouped into four basic categories of systems that determine position on the basis of the following:

- Cell of origin *(nearest cell)*
- Distance *(lateration)*
- Angle *(angulation)*
- Location patterning *(pattern recognition)*

An RTLS system designer can choose to implement one or more of these techniques. This may be clearly seen in some approaches attempting to optimize performance in two or more environments with very different propagation characteristics. An example of this is an RTLS system attempting to yield optimal performance for both indoor and outdoor applications by using two different techniques. It is not unusual to hear arguments supporting the case for a fifth category that encompasses RTLS systems that sense and measure position using a combination of at least two of these methods.

Keep in mind that regardless of the underlying positioning technology, the "real-time" nature of an RTLS is only as real-time as the most current timestamps, signal strengths, or angle-of-incidence measurements. The timing of probe responses, beaconing rates, and location server polling intervals can introduce discrepancies seen between actual and reported device position from one reporting interval to another.

The "Location Tracking Approaches" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations* provides a foundation in the technical aspects of traditional location tracking and positioning systems. To better comprehend the differences between traditional approaches and RF Fingerprinting, this section is highly recommended reading, because it thoroughly explains the concepts of cell of origin, time of arrival (ToA), time difference of arrival (TDoA), angle of arrival (AoA), and location patterning.

What is RF Fingerprinting?

Cisco RF Fingerprinting refers to a new and innovative approach that significantly improves the accuracy and precision available with traditional signal strength lateration techniques. Cisco RF Fingerprinting offers the simplicity of an RSSI-based lateration approach with customized calibration capabilities and improved indoor performance.

RF Fingerprinting significantly enhances RSS lateration through the use of RF propagation models developed from data gathered in the target environment or environments very similar to it. RF Fingerprinting offers the ability to calibrate an RF model to a particular environment in a fashion similar to (but more expeditious than) that of location patterning. But unlike location patterning, a unique custom site calibration is not always required, especially in situations where multiple floors of similar construction, contents, and layout are deployed.

Cisco RF fingerprinting offers several other key advantages over the approaches described in the "Location Tracking Approaches" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations*:

- Uses existing LWAPP-enabled Cisco Unified Networking components—Unlike some other solutions, Cisco LBS with RF Fingerprinting is a 100 percent Wi-Fi RTLS without the need for specialized time-based receivers or other specialized hardware that must be mounted alongside each access point. The Cisco Location Appliance is added to support location and statistics history and serves as a centralized positioning engine for the simultaneous tracking of up to 2500 devices per appliance.
- No proprietary client hardware or software required—The Cisco RF Fingerprinting-based RTLS solution is implemented as a network-side model and not client-side. Because of this, Cisco RF Fingerprinting can provide location tracking for a wide variety of industry-standard Wi-Fi clients (not just WinXP/2000/PPC) *without the need to load proprietary client tracking software or wireless drivers in each client*. This includes popular VoIP handsets such as the Cisco 7920 and other devices for which proprietary location tracking client software is not readily available nor installable.
- Supports popular Wi-Fi active RFID asset tags—Because the Cisco LBS solution implements RF Fingerprinting as a network-side model, there is no dependency on proprietary software being resident in RFID asset tags. This enables the Cisco LBS solution to interoperate with active RFID asset tags from popular vendors including AeroScout and PanGo Networks. Cisco also publishes a complete RFID tag specification that is available to Cisco Technology Partners and encourages the development of interoperable active RFID tag hardware. The Cisco LBS solution is capable of tracking other Wi-Fi active RFID tags that can be configured to authenticate/associate to the underlying installed Cisco centralized WLAN infrastructure as a WLAN client.

- Better accuracy and precision—The Cisco RF Fingerprinting approach yields significantly better performance than solutions employing pure triangulation or signal strength lateration techniques. These techniques typically do not account for effects of attenuation in the environment, making them highly susceptible to reductions in performance. The advantages of Cisco RF Fingerprinting technology start where these traditional approaches leave off. Cisco RF Fingerprinting begins with a significantly better understanding of RF propagation as it relates specifically to the environment in question. Except for the calibration phase in location patterning approaches, none of the traditional lateration or angulation techniques take environmental considerations directly into account in this manner. RF Fingerprinting then goes a step further and applies statistical analysis techniques to the set of collected calibration data. This allows the Cisco Location Appliance to further refine predicted location possibilities for mobile clients, culling out illogical or improbable data and refining accuracy. The net result of these methods is not only better accuracy but significantly improved precision over traditional solutions.
- Reduced calibration effort—Cisco RF Fingerprinting technology offers the key advantages of an indoor location patterning solution but with significantly less effort required for system calibration. While both approaches support on-site calibration, the Cisco RF Fingerprinting approach requires less frequent re-calibration and can operate with larger inter-access point spacing. Cisco RF Fingerprinting can also share RF models among similar types of environments and includes pre-packaged calibration models that can facilitate rapid deployment in typical indoor office environments.

Additional information regarding RF Fingerprinting and all four of these key advantages can be found in the “Location Based Services Architecture” section of *Wi-Fi Location-Based Services: Design and Deployment Considerations*.

Overall Architecture

The overall architecture of Cisco LBS is shown in Figure 15-1.

Figure 15-1 Cisco Location-Based Services Solution Architecture

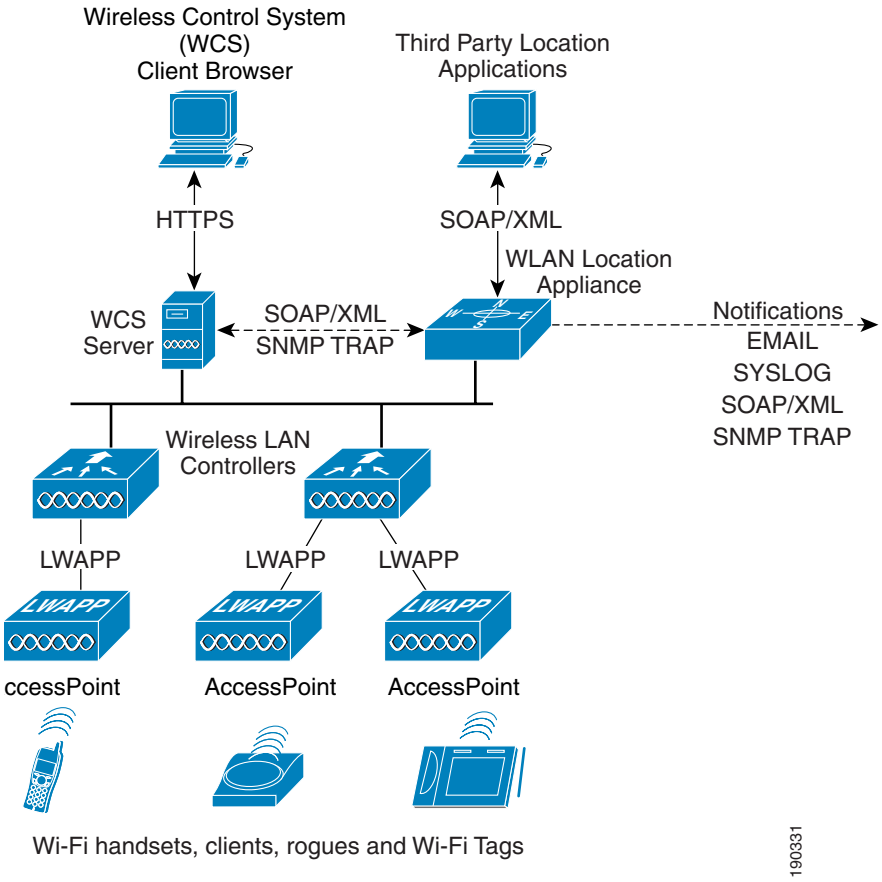

Access points forward received signal strength information to WLAN controllers with regard to the observed signal strength of any Wi-Fi clients, 802.11 active RFID tags, rogue access points, or rogue clients. In normal operation, access points focus their collection activities for this information on their primary channel of operation, going off-channel and scanning the other channels in the regulatory channel set of the access point periodically. The collected information is forwarded to the WLAN controller to which the access point is currently registered. Each controller manages and aggregates all such signal strength information coming from its access points. The location appliance uses SNMP to poll each controller for the latest information regarding each tracked category of devices. In the case of a location tracking system deployed without a location appliance, the Cisco Wireless Control System (WCS) retrieves this information from the appropriate controller(s) directly.

WCS and the location appliance exchange information regarding calibration maps and network designs during a process known as *synchronization*. During a *network design synchronization* between WCS and the location appliance, the "up-to-date" partner updates the design and calibration information of the "out-of-date" partner. The location appliance synchronizes with each controller containing access points participating in location tracking during *controller synchronization*. Synchronization occurs either on-demand or as a scheduled task, the timing of which is determined by the **Administration > Scheduled Tasks** main menu option under the WCS main menu bar.

Location information is displayed to the end user using a *location client* application in conjunction with the Cisco Wireless Location Appliance. Typically this role is fulfilled by the Cisco WCS.

Note For important information regarding compatibility between versions of WCS and the Cisco Wireless Location Appliance, see *Release Notes for Cisco Wireless Location Appliance* at the following URL: http://www.cisco.com/en/US/products/ps6386/prod_release_note09186a00806b5ec7.html.

As described in subsequent sections of this document, the WCS is capable of displaying a wide range of information regarding the location of clients, asset tags, rogue access points, and rogue clients. However, location client functionality is not limited to WCS. Other third-party applications may be written in accordance with the Cisco Location Appliance Application Programming Interface (API) as well. Using the Simple Object Access Protocol (SOAP)/Extensible Markup Language (XML) protocol, these applications can also serve as location clients to the Wireless Location Appliance (see Figure 15-1). The Cisco Location Appliance is also capable of issuing notifications to external systems via e-mail (SMTP), Syslog, SNMP traps, or the SOAP/XML protocol.

Note Additional information regarding the overall architecture of the Cisco LBS solution can be found in the "Location-Based Services Architecture" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations* located at http://www.cisco.com/univercd/cc/td/doc/solution/wifidesi.pdf.

Role of the Cisco Wireless Location Appliance

When a Cisco Location Appliance is added into a Cisco LWAPP-enabled Unified Wireless Network with an appropriately licensed version of WCS, the location appliance assumes responsibility for several important tasks. Key among these are the execution of positioning algorithms, maintenance of calibration information, triggering and dispatch of location notifications, and the ongoing processing of historical location and statistics information. WCS acts in concert with the location appliance by serving as the user interface (UI) for the services the location appliance provides. Although it is possible to access the location appliance directly via SSH or a console session, all end user interaction with the location appliance is typically via WCS or a third-party location client application (except for initial setup of the location appliance and whenever it is necessary to quiesce the appliance).

The integration of a Cisco Location Appliance into a Cisco Unified Wireless Network architecture immediately enables improvements in network location capabilities such as the following:

- Scalability—Adding a Cisco Location Appliance greatly increases the scalability of the Cisco LBS solution from on-demand tracking of a single device to a maximum capacity of 2500 devices (WLAN clients, RFID tags, rogue access points, and rogue clients). For deployments requiring support of greater numbers of devices, additional location appliances can be deployed and managed under a common WCS.
- Historical and statistics trending—The appliance records and maintains historical location and statistics information, which is available for viewing via WCS.
- Location notifications—The Cisco Location Appliance can dispatch location-based event notifications via e-mail, Syslog, SNMP traps, and SOAP/XML directly to specified destinations. These notifications can be triggered simply if the location of a client or asset changes, the client or asset strays beyond set distances from pre-determined marker locations, or the client or asset otherwise becomes missing or enters/leaves coverage areas. Notifications can also be generated for asset tag battery levels (that is, low battery notification).
- SOAP/XML Location Application Programming Interface (API)—The Location Appliance API allows customers and partners to create customized location-based programs that interface with the Cisco Wireless Location Appliance. These programs can be developed to support a variety of unique and innovative applications including real-time location-based data retrieval, telemetric device

management, workflow automation, enhanced WLAN security, and people or device tracking. The API provides a mechanism for inserting, retrieving, updating, and removing data from the Cisco Wireless Location Appliance configuration database using an XML SOAP interface. Developers can access the Cisco Wireless Location Appliance provisioning services using XML and exchange data in XML format. The location appliance API is available and licensable to the Cisco development community along with tools to facilitate solution development. Integration support is available via the Cisco Developer Services Program. For complete details, see http://www.cisco.com/go/developersupport.

Solution Performance

When discussing the performance of a positioning system, the metric that is most familiar and significant is *accuracy,* which typically refers to the quality of the information being received. *Location accuracy* refers specifically to the quantifiable error distance between the estimated location and the actual location of the mobile device.

In most real-world applications, however, a statement of location accuracy has little value without the ability of the solution to repeatedly and reliably perform at this level. *Precision* is a direct measure reflecting on the reproducibility of the stated location accuracy. Any indication of location accuracy should therefore include an indication of the confidence interval or percentage of successful location detection as well, otherwise known as the *location precision*.

When properly deployed, the accuracy and precision of the Cisco LBS solution in indoor deployments is represented in two ways, as follows:

- Accuracy of less than or equal to 10 meters, with 90 percent precision
- Accuracy of less than or equal to 5 meters, with 50 percent precision

In other words, given proper design and deployment of the system, the error distance between the reported device location and the actual location should, in 90 percent of all reporting instances, be within 10 meters or less. In the remaining 10 percent of all reporting instances, the error distance may be expected to exceed 10 meters. Note that these specifications apply only to solutions using RF Fingerprinting; namely, the use of a WCS licensed for location usage (with or without a location appliance).

For applications that require better performance than an accuracy of 10 meters with 90 percent precision, the Cisco LBS solution can deliver accuracy of 5 meters but with 50 percent precision. Or stated another way, in 50 percent of all reporting instances, it can be reasonably expected that the error distance between the reported and the actual location will exceed 5 meters. In addition, the *location inspection* tool (a new feature with release 4.0 of WCS and 2.1 of the location appliance) can display various levels of accuracy and precision from 2m to 100m and the areas of your environment that can meet these accuracy targets. Using the location inspection tool in conjunction with new predictive tools (available in WCS release 4.0 and Location Appliance release 2.1), such as the *location planner* and the *location readiness*, the network designer now can not only plan for achieving stated performance goals but verify that these targets are indeed being met.

What Devices Can Be Tracked

The Cisco LBS solution can provide position tracking information for the following:

- *Standard WLAN clients* or *Wi-Fi 802.11 active RFID tags* that are probing, are associated or attempting association with your controller-based location-aware wireless LAN infrastructure. This includes PanGo Locator LAN RFID tags and other RFID tags that are capable of successfully authenticating and associating to the underlying WLAN infrastructure. These types of wireless LAN clients are displayed on the WCS location floor maps using a blue-rectangular icon.
- *802.11 active RFID asset tags* (which communicate via Layer 2 multicasts and do not associate to the WLAN infrastructure) are displayed on WCS floor maps as a yellow tag icon.
- *Rogue access points* are access points that are detected by the wireless LAN infrastructure and determined not to be members of the same mobility group or WLAN system. These are indicated on WCS location floor maps using a skull-and-crossbones within a black circle.
- *Rogue clients* are clients associated to rogue access points. Rogue clients are displayed on the WCS location floor maps using a black rectangle icon with a skull-and-crossbones.

Note Comprehensive information regarding each trackable class of device that can be displayed by the Cisco LBS solution for each is found in the "Location-Based Services Architecture" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations* located at http://www.cisco.com/univercd/cc/td/doc/solution/wifidesi.pdf.

Installation and Configuration

Installing and Configuring the Location Appliance and WCS

Detailed procedures for installing and configuring the Cisco Wireless Location Appliance and WCS can be found via the references mentioned in the "Installation and Configuration" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations*.

Configuration of the parameters listed under the WCS Location Server > Administration menu are discussed in the document entitled *Cisco Location Appliance Configuration Guide: Editing Location Server Properties* at the following URL: http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a00806b5b10.html.

However, there are additional ramifications associated with making changes to the factory defaults that need to be carefully considered. This and other valuable information that a designer of a location-enabled wireless LAN should consider can be found in the "Installation and Configuration" section in *Wi-Fi Location-Based Services: Design and Deployment Considerations,* including the following:

- History parameters
  - History archive period
  - History data pruning
- Advanced parameters
  - Absent data cleanup interval
  - DB disk memory
  - Run Java GC
  - Defragment database

  - DB free size
- Location parameters
  - Enable calculation time
  - Enable OW location
  - Relative RSSI discard time
  - Absolute RSSI discard time
  - RSSI cutoff
- Location server notification parameters
- Location server dual Ethernet operation
- Location server time synchronization
- Setting passwords for the Wireless Location Appliance
- Proper shutdown (quiescing) of the Wireless Location Appliance

Deployment Best Practices

Location-Aware WLAN Design Considerations

The past decade has witnessed the best practice design of enterprise-ready wireless LANs evolve from being centered around the model of maximum coverage using minimum AP count to a new model where coverage uniformity and proper cell-to-cell overlap are the predominant criteria. This has been driven by increasing interest in deploying new wireless applications such as wireless voice with its intolerance for large amounts of dropped packets and high roaming delays. In a similar fashion, deploying location-based applications using a Wi-Fi wireless LAN requires a modification of the current approach, both in how new "location-aware" installations are designed, and also in how an existing deployment is augmented or retrofitted to take advantage of new location-tracking applications. For location tracking to function optimally, the correct number of access points along with proper access point placement is necessary to assure that mobile devices are properly detected as they move about in the WLAN environment.

The "Deployment Best Practices" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations* discusses in great detail several best-practice recommendations for location-aware WLAN deployments. These best-practice recommendations are briefly described here:

- Minimum detected received signal thresholds—For mobile devices to be tracked properly, it is highly recommended that access points report mobile device RSSI to their respective controllers at levels meeting or exceeding the *RSSI cutoff* value that is configured in WCS. A minimum of three access points (and preferably four or more for optimum accuracy) should be reporting this level of signal strength or better for any device being localized. Mobile device RSSI reported below this level may be discarded by the location appliance.
- Correct access point placement—Proper placement and density of access points is critical to achieving the quoted performance of the Cisco location tracking solution. In many office wireless LANs, access points are distributed throughout interior spaces, thereby providing coverage to the surrounding work areas. These locations are usually selected on the basis of coverage, WLAN bandwidth, channel re-use, cell-to-cell overlap, security, aesthetics, and deployment feasibility. In a location-aware WLAN design, however, access points must not be located based solely on these criteria.

- Correct access point density—Access point density also has a significant effect on location tracking performance. Although there is no single steadfast rule that yields the proper density in every environmental situation, a good suggested starting point is to incorporate the signal threshold and placement suggestions made in the "Deployment Best Practices" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations*. Chief among these is the adherence to an inter-access point separation of 50 to 70 feet, which often results in one location-aware access point being deployed approximately every 2500 to 4900 square feet.
- Minimizing excessive co-channel interference—In many cases, location-based services are added or retrofitted to an existing wireless design, some of which encompass wireless voice handheld devices (such as the Cisco 7920). When designing a location-aware solution that will be used in conjunction with such latency-sensitive application devices, special care needs to be taken to ensure that excessive co-channel interference is not introduced into the environment. The needs of an optimal location-aware design must be carefully balanced against the stringent requirements of a properly designed wireless voice infrastructure.
- Avoiding location display "jitter"—At times, devices appear to move on location displays even though they are known to physically be at rest. This is because of a variety of factors, including movement of other objects in the environment and slight changes in the orientation of the client antenna over time. Release 4.0 of WCS and 2.1 of the Location Appliance introduce the concept of *location smoothing* to assist in counteracting this phenomena and stabilize location jitter for clients that are not in constant motion.
- Multi-floor structures—In multi-floor structures such as office buildings, the location appliance must continually analyze all available signal strength data and make determinations as to which floor each mobile device is currently resident. The location appliance does this by comparing the detected signal strength of the client from access points located on each floor, assigning metrics, and then undergoing a series of calculations to determine the best placement. Understanding the mechanics behind how this is done allows the network designer to lower the potential for floor mis-detects.
- Multi-domain design considerations—In release 2.1 of the Location Appliance, the capacity of the system has increased from 1500 to 2500 total devices, which includes WLAN clients, asset tags, rogue access points, and clients. When combined with the expanded management capacities available using release 4.0 of WCS, a single location appliance and WCS management system should suffice for the majority of applications. However, in larger networks, it may be necessary to use either a single WCS server with multiple location appliances or multiple WCS servers with one or more location appliances, each to address the largest of deployments.
- Antenna considerations—A listing of the supported access point and antenna combinations for use with the Cisco LBS solution, tips on third party antennas, and antenna orientation best practices.
- Site calibration—Important tips on performing site calibrations, calibration validity, choosing a calibration client, and improving overall calibration performance.

Traffic Considerations

The Cisco Wireless Location Appliance and the Cisco WCS are part of the Cisco Unified Wireless Network with each deployed as a separate hardware component for optimum scalability and maximum flexibility. Generally speaking, when all components are deployed in a campus arrangement via a well-designed 10/100/1000 infrastructure wired LAN, bandwidth is typically sufficient for proper operation of the LBS solution. In deployments supporting a large number of geographically distributed locations, further consideration with regard to data traffic load may be required.

The "Deployment Best Practices" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations* provides valuable in-depth discussion and traffic flow analysis of the data flows between the location appliance, WLAN controllers, and WCS.

RFID Tag Considerations

The majority of RFID tags currently produced are *passive* RFID tags, consisting basically of a micro-circuit and an antenna. They are referred to as passive tags because the only time in which they are actively communicating is when they are within the RF field of a passive RFID tag reader or *interrogator*.

Another type of common RFID tag in the current marketplace is known as the *active* RFID tag, which usually contains a battery that directly powers RF communication. This onboard power source allows an active RFID tag to transmit information about itself at great range, either by constantly *beaconing* this information to a RFID tag reader or by transmitting only when it is prompted to do so. Active tags are usually larger in size and can contain substantially more information (because of higher amounts of memory) than do pure passive tag designs.

The "RFID Tag Considerations" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations* provides readers who are new to RFID with a foundation in both active and passive tag technologies. Among other areas, this section comprehensively discusses the following:

- Passive RFID technology—Passive and semi-passive RFID tags
- Active RFID technology—Beaconing, transponder and 802.11 (Wi-Fi) RFID tags
- Using RFID tags with the Location Appliance—Compatible RFID tags, enabling asset tag tracking, configuring asset tags, and using 802.11b tags on 802.11g networks. This section includes a detailed examination of 802.11 active RFID tags from both various suppliers.

The SOAP/XML Application Programming Interface

To facilitate the deployment of location-based applications in the enterprise, the Cisco Wireless Location Appliance is equipped with a rich SOAP/XML API. Applications can make use of the location information contained within the location appliance by importing components via the API such as entire network maps including buildings, floors, access points, coverage areas, and device lists. Actionable data can also be imported, such as recent and historical location as well as statistical device information. Location-based alarms and notifications can be triggered in applications through area boundary definitions, allowed areas and allowed distances. All of these capabilities allow the SOAP/XML API interface to the Cisco Wireless Location Appliance API to be used for integration with external software applications such as E911, asset management, enterprise-resource-planning (ERP) tools, and workflow automation systems that are location-enabled.

From a high-level perspective, a third-party application system can use the SOAP/XML API to participate as a member of a system consisting of the following four basic components:

- Location client—The location client is the recipient of location data that is processed and stored by the location server.
- Control client—The control client administers the location server as well as having the capability to write/read all location data contained on the server.
- Location server—The location server provides the location services for a network or part of a network.

- Wireless LAN system—All the monitored mobile devices (tags, mobile stations, rogue clients, and access points) serving as key components of the wireless network as well as the embedded software contained within WLAN controllers.

The "SOAP/XML Application Programming Interface" section of *Wi-Fi Location-Based Services: Design and Deployment Considerations* describes all four of these basic components in much further detail and briefly examines a Cisco Technology Partner location client implementation.

APPENDIX A

Excerpt of Configuration Audit Exchange, WCS <-> 4400 WLAN Controller

Figure A-1 Configuration Audit Ethernet Trace Excerpt

No. .	Time	Source	Destination	SrcPort	DstPort	Protocol	Bytes	Info
173	0.011	wcswindows	AeS_4402_2	1064	snmp	SNMP	608	GET SNMPv2-SMI::enterprises.14179.2.3.3.1.24.0 SNMPv2-SMI::enterpri
174	0.011	AeS_4402_2	wcswindows	snmp	1064	SNMP	640	RESPONSE SNMPv2-SMI::enterprises.14179.2.3.3.1.24.0 SNMPv2-SMI::ent
175	0.022	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
176	0.023	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
177	0.024	wcswindows	AeS_4402_2	1064	snmp	SNMP	165	GET SNMPv2-SMI::enterprises.14179.2.4.1.1.5.0 SNMPv2-SMI::enterpris
178	0.024	AeS_4402_2	wcswindows	snmp	1064	SNMP	178	RESPONSE SNMPv2-SMI::enterprises.14179.2.4.1.1.5.0 SNMPv2-SMI::ente
179	0.031	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
180	0.031	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
181	0.032	wcswindows	AeS_4402_2	1064	snmp	SNMP	146	GET SNMPv2-SMI::enterprises.14179.2.4.1.6.9.0 SNMPv2-SMI::enterpris
182	0.033	AeS_4402_2	wcswindows	snmp	1064	SNMP	152	RESPONSE SNMPv2-SMI::enterprises.14179.2.4.1.6.9.0 SNMPv2-SMI::ente
183	0.040	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
184	0.040	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
185	0.042	wcswindows	AeS_4402_2	1064	snmp	SNMP	244	GET SNMPv2-SMI::enterprises.14179.2.4.1.6.1.0 SNMPv2-SMI::enterpris
186	0.042	AeS_4402_2	wcswindows	snmp	1064	SNMP	255	RESPONSE SNMPv2-SMI::enterprises.14179.2.4.1.6.1.0 SNMPv2-SMI::ente
187	0.050	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
188	0.050	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
189	0.053	wcswindows	AeS_4402_2	1064	snmp	SNMP	665	GET SNMPv2-SMI::enterprises.14179.2.3.2.1.17.0 SNMPv2-SMI::enterpri
190	0.054	AeS_4402_2	wcswindows	snmp	1064	SNMP	700	RESPONSE SNMPv2-SMI::enterprises.14179.2.3.2.1.17.0 SNMPv2-SMI::ent
191	0.063	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
192	0.063	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
193	0.064	wcswindows	AeS_4402_2	1064	snmp	SNMP	165	GET SNMPv2-SMI::enterprises.14179.2.4.2.1.3.0 SNMPv2-SMI::enterpris
194	0.065	AeS_4402_2	wcswindows	snmp	1064	SNMP	178	RESPONSE SNMPv2-SMI::enterprises.14179.2.4.2.1.3.0 SNMPv2-SMI::ente
195	0.072	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
196	0.072	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
197	0.073	wcswindows	AeS_4402_2	1064	snmp	SNMP	146	GET SNMPv2-SMI::enterprises.14179.2.4.2.6.9.0 SNMPv2-SMI::enterpris
198	0.073	AeS_4402_2	wcswindows	snmp	1064	SNMP	152	RESPONSE SNMPv2-SMI::enterprises.14179.2.4.2.6.9.0 SNMPv2-SMI::ente
199	0.081	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
200	0.081	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
201	0.083	wcswindows	AeS_4402_2	1064	snmp	SNMP	244	GET SNMPv2-SMI::enterprises.14179.2.4.2.6.1.0 SNMPv2-SMI::enterpris
202	0.083	AeS_4402_2	wcswindows	snmp	1064	SNMP	255	RESPONSE SNMPv2-SMI::enterprises.14179.2.4.2.6.1.0 SNMPv2-SMI::ente
203	0.091	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
204	0.091	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
205	0.092	wcswindows	AeS_4402_2	1064	snmp	SNMP	106	GET SNMPv2-SMI::enterprises.14179.2.3.1.29.0 SNMPv2-SMI::enterprise
206	0.092	AeS_4402_2	wcswindows	snmp	1064	SNMP	108	RESPONSE SNMPv2-SMI::enterprises.14179.2.3.1.29.0 SNMPv2-SMI::enter
207	0.102	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
208	0.102	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
209	0.103	wcswindows	AeS_4402_2	1064	snmp	SNMP	106	GETBULK SNMPv2-SMI::enterprises.14179.2.5.22.1.2 SNMPv2-SMI::enterp
210	0.103	AeS_4402_2	wcswindows	snmp	1064	SNMP	885	RESPONSE SNMPv2-SMI::enterprises.14179.2.5.22.1.2.12.48.48.48.98.10
211	0.118	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
212	0.119	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
213	0.119	wcswindows	AeS_4402_2	1064	snmp	SNMP	108	GETBULK SNMPv2-SMI::enterprises.14179.1.2.3.23.1.2 SNMPv2-SMI::ente
214	0.120	AeS_4402_2	wcswindows	snmp	1064	SNMP	488	RESPONSE SNMPv2-SMI::enterprises.14179.1.2.3.24.0 SNMPv2-SMI::enter
215	0.124	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
216	0.124	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0
217	0.125	wcswindows	AeS_4402_2	1064	snmp	SNMP	88	GETBULK SNMPv2-SMI::enterprises.14179.2.5.6.1.2
218	0.126	AeS_4402_2	wcswindows	snmp	1064	SNMP	287	RESPONSE SNMPv2-SMI::enterprises.14179.2.5.11.1.1.5.97.100.109.105.
219	0.130	wcswindows	AeS_4402_2	1064	snmp	SNMP	84	GET SNMPv2-MIB::sysUpTime.0
220	0.130	AeS_4402_2	wcswindows	snmp	1064	SNMP	87	RESPONSE SNMPv2-MIB::sysUpTime.0

190798

WCS Event and Alarm Severities

Double-quotations enclose variables that are replaced with resource names when the message is displayed.

Critical Events and Alarms

1. The PoE controller has failed on the controller “{0}”.
2. AP “{0}”, interface “{1}” is down on controller “{2}”.
3. AP “{0}” disassociated from controller “{1}”
4. Controller “{0}”. RADIUS server(s) are not responding to authentication requests.
5. Port “{0}” is down on controller “{1}”.
6. Rogue AP “{0}” is on wired network.
7. User “{1}” with IP address “{0}” has made too many unsuccessful login attempts.
8. AP “{0}” with protocol “{1}” on controller “{2}” is contained as a rogue, preventing service.
9. Fake AP or other attack may be in progress. Rogue AP count on system “{0}” has exceeded the security warning threshold of “{1}”.
10. Fake AP or other attack may be in progress. Rogue AP count on AP with MAC address “{0}” associated with controller “{2}” has exceeded the security warning threshold of “{1}”.
11. Controller “{0}” detected duplicate IP address “{0}” being used by machine with MAC address “{1}”.
12. AP “{0}” on controller “{3}” detected duplicate IP address “{2}” being used by machine with MAC address “{1}”.
13. The AP “{0}” with protocol “{1}” received a message with a large NAV field and all traffic on the channel has been suspended. This is most likely a malicious DoS attack.
14. The AP “{1}” received a WPA MIC error on protocol “{2}” from Station “{0}”. Counter measures have been activated and traffic has been suspended for 60 seconds.
15. Controller “{0}” is unreachable.
16. IDS signature attack detected on controller “{0}”. The signature type is “{1}”, signature name is “{2}”, and signature description is “{3}”.
17. Transmitter failure detected on the “{0}” radio of AP “{1}” on controller “{2}”.
18. Receiver failure detected on the “{0}” radio of AP “{1}” on controller “{2}”.

19. AP impersonation with MAC “{0}” is detected by authenticated AP “{1}” on “{2}” radio and Slot ID “{3}”.
20. AP functionality has been disabled for key “{0}”, reason being “{1}” for feature set “{2}”.
21. AP “{1}” is unable to associate. The regulatory domain configured on it “{3}” does not match the controller “{0}” country code “{2}”.
22. CPU Receive Multicast Queue is full on controller “{0}”.
23. Failed to authorize AP “{0}”. Authorization entry does not exist in AP authorization list of controller “{1}” .
24. Failed to authorize AP “{0}”. AP’s authorization key does not match with SHA1 key in AP authorization list of controller “{1}” .
25. Failed to authorize AP “{0}”. Controller “{1}” could not verify the self-signed certificate from the AP.
26. Failed to authorize AP “{0}”. AP has a self-signed certificate, whereas the AP authorization list of controller “{1}” has manufactured installed certificate for this AP.
27. Radio with MAC address “{0}” and protocol “{1}” is down. The reason is “{2}”.
28. Radio with MAC address “{0}” and protocol “{1}” that has joined controller “{2}” has invalid interface. The reason is “{3}”.
29. The Cisco Intrusion Detection System “{0}” has detected a possible intrusion attack by the wireless client “{1}”.
30. Controller “{0}” is “{1}” with the central time server.
31. MFP configuration of the WLAN was violated by the radio interface “{0}” and detected by the radio interface “{1}” of the AP with MAC address “{2}”. The violation was “{3}”.
32. Guest user “{1}” deleted on controller “{0}”.

Major Events and Alarms

1. The radios associated with controller “{0}” exceeded license count “{1}”. The current number of radios on this controller is “{2}”.
2. The sensed temperature on the controller “{0}” is too high. The current sensed temperature is “{1}”.
3. The sensed temperature on the controller “{0}” is too low. The current sensed temperature is “{1}”.
4. The temperature sensor failed on the controller “{0}”. Temperature is unknown.
5. Adhoc rogue “{0}” was found and has been auto-contained as per WPS policy.
6. Rogue AP “{0}” was advertising the SSID and has been auto-contained as per WPS policy.
7. Trusted AP “{0}” has invalid encryption. It is using “{1}” instead of “{2}”. It has been auto-contained as per WPS policy.
8. Trusted AP “{0}” has invalid radio policy. It is using “{1}” instead of “{2}”. It has been auto-contained as per WPS policy.
9. Trusted AP “{0}” has invalid SSID. It has been auto-contained as per WPS policy.
10. Trusted AP “{0}” is missing or has failed.
11. Trusted AP “{0}” on controller “{3}” has invalid preamble. It is using “{1}” instead of “{2}”. It has been auto-contained as per WPS policy.

12. Keepalive messages are lost between master and controller “{0}”.

Minor Events and Alarms

1. AP “{0}”, interface “{1}”. Load threshold violated.
2. AP “{0}”, interface “{1}”. Noise threshold violated.
3. AP “{0}”, interface “{1}”. Interference threshold violated.
4. AP “{0}”, interface “{1}”. Coverage threshold of “{3}” is violated. Total number of clients is “{5}” and number failed clients is “{4}”.
5. Controller “{0}”. User authentication from controller “{0}” failed for user name “{1}” and user type “{2}”.
6. Client “{0}”, which was associated with AP “{1}”, interface “{2}” is excluded. The reason code is “{3}”.
7. IPsec IKE negotiation failure from remote IP address “{0}”.
8. IPsec invalid cookie from remote IP address “{0}”.
9. Rogue AP “{0}” with SSID “{3}” and channel number “{4}” is detected by AP “{1}”. Radio type “{2}” with RSSI “{5}” and SNR “{6}”.
10. The WEP key configured at the station may be wrong. Station MAC address is “{0}”, AP MAC is “{1}”, and Slot ID is “{2}”.
11. AP “{0}” with static IP configured as “{2}” has fallen back to the working DHCP address “{1}”.
12. Absence of <Element> with MAC <macAddress>, last seen at <timestamp>
13. <Element> with MAC <macAddress> is <In | Out> the Area <campus | building | floor | coverageArea>
14. <Element> with MAC <macAddress> has moved beyond <specifiedDistance> ft. of marker <MarkerName>, located at a range of <foundDistance> ft.

Clear Events and Alarms

1. AP “{0}”, interface “{1}” is up.
2. AP “{0}”, interface “{1}”. Load changed to acceptable.
3. AP “{0}”, interface “{1}”. Noise changed to acceptable.
4. AP “{0}”, interface “{1}”. Interference changed to acceptable.
5. AP “{0}”, interface “{1}”. Coverage changed to acceptable.
6. Port “{0}” is up on controller “{1}”.
7. Rogue AP “{0}” is removed; it was detected as rogue AP by AP “{1}”. Radio type “{2}”.
8. Rogue AP “{0}” is not able to connect to the wired network.
9. The temperature sensor is working now on the controller “{0}”. The sensed temperature is “{1}”.
10. Controller “{0}” is reachable.
11. Adhoc rogue “{0}” was found and was auto-contained. The alert state is clear now.
12. Rogue AP “{0}” was advertising the SSID and was auto-contained. The alert state is clear now.

13. Trusted AP "{0}" had invalid encryption. The alert state is clear now.
14. Trusted AP "{0}" had invalid radio policy. The alert state is clear now.
15. Trusted AP "{0}" had invalid SSID. The alert state is clear now.
16. Trusted AP "{0}" is missing or has failed. The alert state is clear now.
17. Controller "{0}" is cleared from IDS signature attack. The wireless system is no longer detecting the intrusion.
18. Transmitter failure cleared on the "{0}" radio of AP "{1}" on controller "{2}".
19. Receiver failure cleared on the "{0}" radio of AP "{1}" on controller "{2}".
20. Trusted AP "{0}" on controller "{3}" had invalid preamble. The alert state is clear now.
21. Radio with MAC address "{0}" and protocol "{1}" is up. The reason is "{2}".
22. Radar has been cleared on channel "{1}", which was detected by AP base radio MAC "{0}" on radio 802.11a.
23. The Cisco Intrusion Detection System "{0}" has cleared the wireless client "{1}" from possibly having generated an intrusion attack.

Informational Events and Alarms

1. Controller "{0}". Configuration saved in flash.
2. Controller "{0}". Multiple users logged in.
3. Controller "{0}". Cold start.
4. AP "{0}" associated with controller "{2}" on port number "{1}".
5. AP "{0}", interface "{1}". Transmit power level changed to "{2}".
6. AP "{0}", interface "{1}". Channel changed to "{2}". Interference energy before update was "{3}" and after update is "{4}".
7. RRM 802.11a grouping done; the new group leader MAC address is "{0}".
8. RRM 802.11b/g grouping done; the new group leader MAC address is "{0}".
9. Controller "{0}". Authentication failure reported
10. Client "{0}" is associated with AP "{1}", interface "{2}".
11. Client "{0}" with user name "{3}" is authenticated with AP "{1}", interface "{2}".
12. Client "{0}" is disassociated from AP "{1}", interface "{2}" with reason code "{3}".
13. Client "{0}" is deauthenticated from AP "{1}", interface "{2}" with reason code "{3}".
14. Client "{0}" has failed authenticating with AP "{1}", interface "{2}". The reason code is "{3}".
15. Client "{0}" failed to associate with AP "{1}", interface "{2}". The reason code is "{3}".
16. Rogue AP "{0}" is cleared explicitly. It is not detected anymore.
17. Fake AP or other attack is cleared now. Rogue AP count on system "{0}" is within the threshold of "{1}".
18. Fake AP or other attack on AP with MAC address "{0}" associated with controller "{2}" is cleared now. Rogue AP count is within the threshold of "{1}".
19. Global "{1}" network status disabled on controller with IP address "{0}".

20. Global “{1}” network status enabled on controller with IP address “{0}”.
21. Radio with MAC address “{0}” and protocol “{1}” has core dump on controller “{2}”.
22. AP “{0}” tried to join controller “{1}” and failed. The controller does not support this kind of AP.
23. Radar has been detected on channel “{1}” by AP base radio MAC “{0}” on radio 802.11a.

APPENDIX C

Example of Wireless LAN Controller Initial Setup

```
Use the '-' character to backup
System Name [Cisco_40:3d:c3]: AeS_4402_1
Enter Administrative User Name (24 characters max): admin
Enter Administrative Password (24 characters max): *****

Service Interface IP Address Configuration [none][DHCP]: none
Service Interface IP Address: 192.168.0.40
Service Interface Netmask: 255.255.255.0

Enable Link Aggregation (LAG) [yes][NO]: NO

Management Interface IP Address: 10.1.56.16
Management Interface Netmask: 255.255.252.0
Management Interface Default Router: 10.1.56.2
Management Interface VLAN Identifier (0 = untagged): 0
Management Interface Port Num [1 to 2]: 1
Management Interface DHCP Server IP Address: 10.1.56.1

AP Transport Mode [layer2][LAYER3]: LAYER3
AP Manager Interface IP Address: 10.1.56.17

AP-Manager is on Management subnet, using same values
AP Manager Interface DHCP Server (10.1.56.1): 10.1.56.1

Virtual Gateway IP Address: 1.1.1.1

Mobility/RF Group Name: Mobility Group 1

Network Name (SSID): testuser
Allow Static IP Addresses [YES][no]: YES

Configure a RADIUS Server now? [YES][no]: no

Enter Country Code (enter 'help' for a list of countries) [US]: US

Enable 802.11b Network [YES][no]: YES
Enable 802.11a Network [YES][no]: YES
Enable 802.11g Network [YES][no]: YES
Enable Auto-RF [YES][no]: YES

Configuration saved!
Resetting system with new configuration...
```

Examples of SNMP Traps

Figure D-1 AP Disassociated, AP Interface Down, and Link Down Traps

Figure D-2 Client Authentication Failure, Client De-Authenticated, and WLC Configuration Saved Traps

Figure D-3 Rogue AP Detected, Rogue AP Removed, and Interference Profile Failed Traps

Figure D-4 AP Current Channel Changed, AP Current Tx Power Changed, AP Load Profile Failed Traps

Sample Monitor > Devices > Access Points Reports

Figure E-1 Load, TX Power, Noise, and Interference Report Options

Access Points > Load

AP Name	Radio	Attached Client Count		Channel Utilization		Receive Utilization	Transmit Utilization	Status
		Actual	Threshold	Actual	Threshold			
sjc14-42b-ap7	802.11b/g	4	12	33%	80%	28%	8%	Okay
sjc14-41b-ap6	802.11b/g	2	12	32%	80%	15%	4%	Okay
sjc14-42b-ap6	802.11b/g	2	12	23%	80%	23%	7%	Okay
sjc14-42b-ap2	802.11b/g	8	12	27%	80%	29%	9%	Okay
sjc14-41b-ap3	802.11b/g	6	12	24%	80%	17%	5%	Okay

Access Points > Dynamic Power Control

AP Name	Radio	Current Power Level	Power Assignment Mode	Recommended Power Level
sjc14-42b-ap7	802.11b/g	2	Automatic	2
sjc14-41b-ap6	802.11b/g	1	Automatic	1
sjc14-42b-ap6	802.11b/g	1	Automatic	1
sjc14-42b-ap2	802.11b/g	1	Automatic	1
sjc14-41b-ap3	802.11b/g	1	Automatic	1

Noise

Per Channel Noise

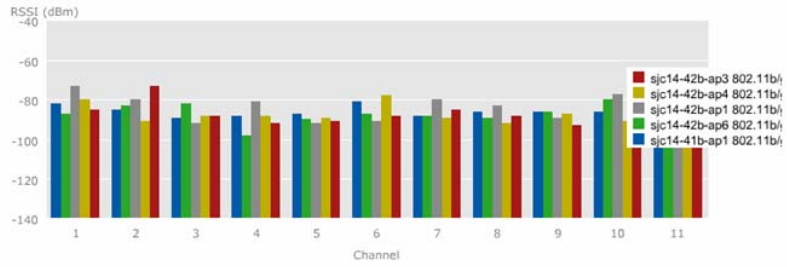

Per Channel Interference

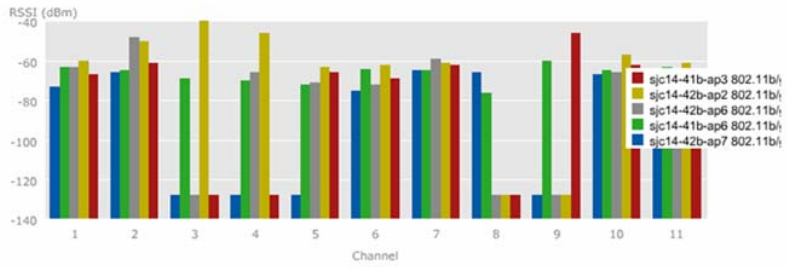

High interference: -40 to 0 dBm
Marginal interference: -100 to -40 dBm
Low interference: -140 to -100 dBm

190804

Figure E-2 Client Distribution by RSSI and SNR, Access Point Up Time