



Wi-Fi Location-Based Services 4.1 Design Guide

May 20, 2008

Americas 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: 408 527-0883

Text Part Number: OL-11612-01

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCVP, the Cisco Logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0612R)

Wi-Fi Location-Based Services 4.1 Design Guide © 2008 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Overview 1-1

- Introduction 1-1
- About the Guide 1-2
 - Target Audience 1-2
 - Objective 1-2
 - Additional Reference Documents 1-3
 - Hardware and Software Components 1-4

CHAPTER 2

Location Tracking Approaches 2-1

- Cell of Origin 2-2
- Distance-Based (Lateration) Techniques 2-3
 - Time of Arrival 2-3
 - Time Difference of Arrival (TDoA) 2-5
 - Received Signal Strength (RSS) 2-7
- Angle-Based (Angulation) Techniques 2-9
 - Angle of Arrival (AoA) 2-9
- Location Patterning (Pattern Recognition) Techniques 2-11
 - Calibration Phase 2-11
 - Operational Phase 2-12

CHAPTER 3

Cisco Location-Based Services Architecture 3-1

- RF Fingerprinting 3-1
- Location-Aware Cisco UWN Architecture 3-4
- Role of the Location Appliance 3-7
 - Location Tracking without a Location Appliance 3-9
- Accuracy and Precision 3-11
- Tracking Clients, Assets and Rogue Devices 3-13
 - Client Probing 3-13
 - Cisco Compatible Extensions Location Measurements 3-14
 - WLAN Clients 3-18
 - 802.11 Active RFID Tags 3-23
 - Rogue Access Points 3-29

- Rogue Clients 3-32
- Workgroup Bridges 3-35
- Cisco Location Control Protocol (LOCP) 3-36
 - Asset Tag Telemetry Using LOCP 3-38
 - Asset Tag Notifications Using LOCP 3-41

CHAPTER 4

Installation and Configuration 4-1

- Installing and Configuring the Location Appliance 4-1
- Configuring Cisco WCS for Location Tracking 4-1
- Configuring Location Appliance History Parameters 4-2
 - History Archive Period 4-2
 - History Database Pruning 4-3
- Configuring Location Appliance Advanced Parameters 4-4
 - Absent Data Cleanup Interval 4-4
 - Memory Information 4-6
 - Advanced Commands 4-6
- Configuring Location Appliance Location Parameters 4-7
 - Enable Calculation Time 4-7
 - Enable OW (Outer Wall) Location 4-8
 - RSSI Discard Times 4-8
 - RSSI Cutoff 4-8
- Configuring Location Appliance Notification Parameters 4-9
 - Queue Limit 4-9
 - Retry Count 4-9
 - Refresh Time 4-10
 - Notifications Dropped 4-10
- Configuring Location Appliance LOCP Parameters 4-10
 - Location Appliance Dual Ethernet Operation 4-10
- Changing Location Appliance Default Passwords 4-11
 - Changing the “root” User Linux System Password 4-11
 - Changing the “admin” Location Server Application Password 4-12
- Location Appliance Time Synchronization 4-14
- Quiescing the Location Appliance 4-15

CHAPTER 5

Best Practices—Location-Aware WLAN Design Considerations 5-1

- Minimum Signal Level Thresholds 5-2
- Access Point Placement 5-5
- Access Point Separation 5-12

Determining Location Readiness	5-18
Location, Voice and Data Coexistence	5-20
Avoiding Location Display Jitter	5-31
Multiple Location Appliance Designs	5-32
Single Management Domain with Multiple Location Domains	5-36
Multiple Management Domains with Multiple Location Domains	5-40
Antenna Considerations	5-44
Third-Party Antennas	5-44
Antenna Orientation and Access Point Placement	5-46
Calibration	5-48
Calibration Validity	5-56
Tips for Successful Calibrations	5-57
Data Collection	5-57
Calibrating Under Representative Conditions	5-58
Recommended Calibration Clients and Techniques	5-59
Calibration of Non-Uniform Environments	5-63
Inspecting Location Quality	5-64
Using Test Points to Verify Accuracy	5-68

CHAPTER 6

RFID Tag Considerations	6-1
RFID Tag Technology	6-1
Passive RFID Tags	6-2
Semi-Passive RFID Tags	6-5
Active RFID Tags	6-6
Beaconing Active RFID Tags	6-7
802.11 Active RFID Tags	6-7
Multimode RFID Tags	6-8
Chokepoint Triggers	6-9
Using Wi-Fi RFID Tags with the Cisco UWN	6-15
Compatible RFID Tags	6-15
Using 802.11b Tags in an 802.11g Environment	6-16
Enabling Asset Tag Tracking	6-17
Enable Asset Tag RF Data Timeout	6-17
Enable Asset Tag Polling	6-18
Enable Asset Tag Display	6-20
Configuring Asset Tags	6-20
Tag Telemetry and Notification Considerations	6-27
Deploying Tag Telemetry	6-27

- Deploying Tag High-Priority Notifications 6-30
- Configuring Tags for Telemetry and Notifications 6-31
- Chokepoint Considerations 6-31
 - Configuring Chokepoint Triggers 6-31
 - Defining Chokepoint Triggers to the Cisco UWN 6-33
 - Chokepoint Trigger Traffic Considerations 6-34

CHAPTER 7

Caveats 7-1

- CSCse14724—Degraded Location Accuracy with Monitor Mode APs 7-1
- CSCsh88795—CCX S36 Beacon Measurement Request Dual-Band Support 7-1
- CSCsi95122—WCS Does Not Dispatch Northbound Emails for Location Notifications 7-2

APPENDIX A

Determining Approximate Roots using Maxima A-1

APPENDIX B

Verifying Detection of Asset Tags in WLAN Controllers B-1

- Asset Tags Detection B-1
- Asset Tags Not Detected B-6
- Verifying Asset Tag Telemetry and Events B-8



CHAPTER 1

Overview

Introduction

802.11 wireless has truly blossomed in the past decade, moving from a technology that was once thought of as primarily a productivity enhancement for vertical industries to one now pervasive throughout society. The wide-spread acceptance of Wi-Fi networks has fueled this dramatic adoption, from deployments in offices and distribution centers to homes and ever-multiplying wireless metropolitan areas. Maturing rapidly and reaching critical mass, this widespread adoption has driven down the cost of wireless infrastructure dramatically and has resulted in the availability of higher quality equipment at lower cost.

The rapid increase in the adoption rate of Wi-Fi coupled with the availability of high quality infrastructure at reasonable cost are key factors behind the flurry of activity regarding Wi-Fi location-based services. Not to be confused with solutions requiring a dedicated, independent infrastructure of location receivers and RFID tag readers, research and development in Wi-Fi location prediction techniques has facilitated the emergence of indoor RF location tracking systems based fundamentally on IEEE 802.11 infrastructure. In combination with the frenetic race to implement RFID systems in the consumer and distribution supply chains, these have all combined to form a “perfect storm” of sorts, transforming what was once a general market passing interest in location-based services into one that well positions 802.11-based location-based services as a potential *must-have* application for Wi-Fi wireless.

It is not difficult to understand why this is so. With integrated location tracking, enterprise wireless LANs become much more valuable as a corporate business asset. This is especially true in today’s fast-paced and highly competitive marketplace, where an otherwise well-positioned enterprise may falter against its peers not because of a lack of necessary assets, but rather due to its inability to quickly locate and re-deploy those assets to address today’s rapidly changing business climate. Enterprise network administrators, security personnel, users, asset owners and others have expressed great interest in location-based services to allow them to better address key issues in their environments, such as the following:

- The need to quickly and efficiently locate 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 rapid location of critical service-impacting assets.
- Improving WLAN planning and tuning capabilities.
- Coordinating Wi-Fi device location with security policy enforcement.
- Determining the location of rogue devices.

- Monitoring the health and status of key assets in their environment and receiving prompt notification of changes.
- Receiving prompt notification when unauthorized addition or removal of assets occurs.

This guide discusses the location-aware Cisco Unified Wireless Network (UWN). It is focused on indoor location-based services design considerations and select deployment topics. References to applicable existing documentation are made throughout the document, and a wealth of material is provided addressing topics such as:

- The fundamentals of positioning technologies including lateration, angulation, and pattern recognition approaches.
- How Cisco RF Fingerprinting operates and how it compares to other approaches.
- The architecture of the location-aware Cisco UWN.
- Design best practices, including voice, data, and location-based service coexistence.
- Tips on proper installation and configuration.

About the Guide

Target Audience

This guide is intended for individuals interested in designing and deploying indoor Cisco wireless LAN (WLAN) solutions that include the Cisco Wireless Location Appliance, the Cisco Wireless Control System (WCS), and other components of the Cisco Unified Wireless Network (UWN).

Objective

This guide is intended to accomplish the following objectives:

- Provide the reader unfamiliar with location-based services with a basic foundation in technical aspects of location tracking and positioning systems. [Chapter 2, “Location Tracking Approaches,”](#) provides substantial background information on positioning system techniques such as cell of origin, time of arrival, time difference of arrival, angle of arrival, and pattern recognition.
- Describe and define RF Fingerprinting, the technology at the heart of the location-aware Cisco UWN. [Chapter 3, “Cisco Location-Based Services Architecture,”](#) discusses the similarities and differences between RF Fingerprinting and other approaches described in [Chapter 2, “Location Tracking Approaches,”](#) and how RF Fingerprinting addresses the deployment of cost-effective indoor Wi-Fi location tracking solutions. This knowledge is useful when comparing the location-aware Cisco Unified Wireless Network to other approaches for indoor location tracking.
- Review the procedures required to install and configure a location-aware Cisco UWN consisting of LWAPP-enabled access points, third-party chokepoint triggers, WLAN controllers, WCS, and the location appliance.
- Provide information that aids in proper installation and understanding of some of the more advanced parameters used (see [Chapter 4, “Installation and Configuration”](#)).
- Describe best practices that should be followed in designing and deploying location-aware wireless LANs. [Chapter 5, “Best Practices—Location-Aware WLAN Design Considerations,”](#) focuses on a variety of topics from access point placement and separation, multiple location appliance designs

and antenna considerations to calibration, and challenging location environments. All the information contained in this section is aimed at assisting designers in optimizing location-aware designs for improved location fidelity.

- Provide the reader having limited exposure to RFID tag technology with a basic understanding of how these various types of tags relate to the location-aware Cisco UWN. [Chapter 6, “RFID Tag Considerations,”](#) provides details regarding RFID asset tags and how these products function. This section also places considerable emphasis on the proper configuration of Cisco WLAN controllers, the WCS, and the location appliance when using RFID tags.

Additional Reference Documents

It is assumed the reader is familiar with the following documents:

- Cisco Wireless Location Appliance Support Documentation for Release 3.0
http://www.cisco.com/en/US/products/ps6386/tsd_products_support_series_home.html
- Cisco Wireless Control System Support Documentation for Release 4.1
http://www.cisco.com/en/US/products/ps6305/tsd_products_support_series_home.html
- Cisco 4400 Series WLAN Controller Support Documentation for Release 4.1
http://www.cisco.com/en/US/products/ps6366/tsd_products_support_series_home.html
- Cisco 2100 Series WLAN Controller Support Documentation for Release 4.1
http://www.cisco.com/en/US/products/ps7206/tsd_products_support_model_home.html
- Cisco Catalyst 3750 Series Integrated Wireless LAN Controllers Support Documentation
http://www.cisco.com/en/US/products/ps6915/tsd_products_support_model_home.html
- Cisco Wireless LAN Controller Module Support Documentation
http://www.cisco.com/en/US/products/ps6730/tsd_products_support_model_home.html
- Cisco Catalyst 6500 Series Wireless Services Module (WiSM) Support Documentation
http://www.cisco.com/en/US/products/ps6526/tsd_products_support_model_home.html



Note

Despite the difference in nomenclature, software Release 3.0 of the Cisco Location Appliance is generally included in any reference made to software Release 4.1 of the Cisco Unified Wireless Network (UWN) within this document.

Additional design considerations surrounding the use of the InnerWireless (formerly PanGo) Vision Locator location client in an integrated Cisco – InnerWireless solution can be found in the following document:

- *Design Considerations for Cisco – PanGo Asset Tracking*
<http://www.cisco.com/univercd/cc/td/doc/solution/pangoex.pdf>

The following guide is recommended as a design reference when considering the deployment of voice over WLAN (VoWLAN) handsets and supporting infrastructure in conjunction with location based services:

- *Voice over Wireless LAN 4.1 Design Guide*
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>

Hardware and Software Components

Table 1-1 lists the hardware and software used in the writing of this guide.


Note

Other supported hardware or software can be found by referring to the information located at the following URL: <http://www.cisco.com/en/US/products/ps6386/index.html>.

Table 1-1 **Tested Hardware and Software**

Location Appliance	
AIR-LOC2700-L-K9 ¹	Location Appliance 2700 Series; software release 3.0.42.0.
Wireless Control System (WCS)	
WCS-STANDARD-K9-4.1.91.0.exe	Wireless Control System Release 4.1.91.0 for Windows 2003 Server ²
WLAN Controllers	
AIR-WLC4402-12-K9	4400 Series WLAN Controller; Release 4.1.185.0
AIR-WLC2106-K9	2106 Series WLAN Controller, Release 4.1.185.0
Access Points	
AIR-LAP1242AG-A-K9	802.11ag LWAPP AP North American; version 12.3(11)JX
External Antennas	
AIR-ANT4941	2.4 GHz, 2.2 dBi Dipole
AIR-ANT5135D-R	5 GHz 3.5 dBi Dipole

1. The Cisco Wireless Location Appliance 2710 (AIR-LOC2710-L-K9) model is the successor to the 2700 (AIR-LOC2700-L-K9) model. There is no functional difference between the 2700 and 2710 models, both models support the same features and functionality.

2. Requires appropriate licensing for Location-Based Services support and total number of access points supported.

Table 1-2 **Additional Hardware and Software Components**

AeroScout Ltd.	System Manager	Version 3.2.20.1
	Network Exciter Manager (ANEM)	version 1.2
	EX-3100 Exciter Manager	1.0.20
	Exciter EX-2000	DSP 216, SB 50008, HW v2
	Exciter EX-3100	310.01
	Exciter EX-3200	DSP 30007, SB 60007, HW v2
	ADP-030 power supply	NA
	ADP-040 power supply	NA
	Tag Manager	3.0.4

Table 1-2 Additional Hardware and Software Components

	Tag Activator BWH1000-02-TA	DSP 22059, FPGA 5201, SB 175, CPLD 5107
	BWH-3000CT T2 RFID Tag	4.33
	EDK-200 T2 Telemetry Tag	4.33
	ExciterConfig.exe	1.3.1



CHAPTER 2

Location Tracking Approaches

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 designer can choose to implement one or more of these techniques. This may be clearly seen in some approaches that attempt to optimize performance in two or more environments with very different propagation characteristics. The popularity of this approach is such that it is often not unusual to hear arguments supporting the case for a fifth category that encompasses RTLS offerings 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 its most current timestamps, signal strength readings, or angle-of-incidence measurements. The timing of probe responses, tag transmissions, and location server polling intervals can introduce discrepancies between the actual and reported device position observed during each reporting interval.

Cell of Origin

One of the simplest mechanisms of estimating approximate location in any system based on RF “cells” is the concept of cell-of-origin (or “associated access point” in Wi-Fi 802.11 systems), as shown in [Figure 2-1](#).


Figure 2-1 Cell of Origin



In its simplest form, this technique makes no explicit attempt to resolve the position of the mobile device beyond indicating the cell with which the mobile device is (or has been) registered. When applied to 802.11 systems, this technique tracks the cell to which a mobile device associates. The primary advantage of this technique is ease of implementation. Cell of origin does not require the implementation of complicated algorithms and thus positioning performance is very fast. Almost all cell-based WLANs and other cellular-based RF systems can be easily and cost-effectively adapted to provide cell of origin positioning capability. However, the overwhelming drawback of pure cell of origin positioning approaches continues to be coarse granularity. For various reasons, mobile devices can be associated to cells that are not in close physical proximity, despite the fact that other nearby cells would be better candidates. This coarse granularity can be especially frustrating when attempting to resolve the actual location of a mobile device in a multi-story structure where there is considerable floor-to-floor cell overlap.

To better determine which areas of the cell possess the highest probability of containing the mobile device, some additional method of resolving location within the cell is usually required. This can either be a manual method (such as a human searching the entire cell for the device) or a computer-assisted method. When receiving cells provide *received signal strength indication (RSSI)* for mobile devices, the use of the *highest signal strength* technique can improve location granularity over the cell of origin. In this approach, the localization of the mobile device is performed based on the cell that detects the mobile device with the highest signal strength. This is shown in [Figure 2-2](#), where the blue rectangular client device icon is placed nearest the cell that has detected it with the highest signal strength.

Figure 2-2 Highest Signal Strength Technique



Using this technique, the probability of selecting the true “nearest cell” is increased over that seen with pure cell of origin. Depending on the accuracy requirements of the underlying business application, performance may be more than sufficient for casual location of mobile clients using the highest signal strength technique. For instance, users intending to use location-based services only when necessary to help them find misplaced client devices in non-mission critical situations may be very comfortable with the combination of price and performance afforded by solutions using the highest signal strength approach. However, users requiring more precise location would find the inability of the highest signal strength technique to isolate the location of a mobile device with finer granularity than that of an entire coverage cell to be a serious limitation. These users are better served by those approaches using the techniques of lateration, angulation, and location patterning that provide finer resolution and improved accuracy. These techniques are discussed in subsequent sections.

Distance-Based (Lateration) Techniques

Time of Arrival

Time of Arrival (ToA) systems are based on the precise measurement of the arrival time of a signal transmitted from a mobile device to several receiving sensors. Because signals travel with a known velocity (approximately the speed of light (c) or ~ 300 meters per microsecond), the distance between the mobile device and each receiving sensor can be determined from the elapsed propagation time of the

signal traveling between them. The ToA technique requires very precise knowledge of the transmission start time(s), and must ensure that all receiving sensors as well as the mobile device are accurately synchronized with a precise time source.

From knowledge of both propagation speed and measured time, it is possible to calculate the distance (D) between the mobile device and the receiving station:

$$D = c (t)$$


where:

- D = distance (meters)
- c = propagation speed of ~ 300 meters / microsecond
- t = time in microseconds

With distance used as a radius, a circular representation of the area around the receiving sensor can be constructed for which the location of the mobile device is highly probable. ToA information from two sensors resolves a mobile device position to two equally probable points. ToA *tri-lateration* makes use of three sensors to allow the mobile device location to be resolved with improved accuracy.

Figure 2-3 illustrates the concept of ToA tri-lateration. The amount of time required for a message transmitted from station X to arrive at receiving sensors A, B, and C is precisely measured as t_A , t_B , and t_C . Given a known propagation velocity (stated as c), the mobile device distance from each of these three receiving sensors can then be calculated as D_A , D_B , and D_C , respectively. Each calculated distance value is used to construct a circular plot around the respective receiving sensor. From the individual perspective of each receiver, station X is believed to reside somewhere along this plot. The intersection of the three circular plots resolves the location of station X as illustrated in Figure 2-3. In some cases, there may be more than one possible solution for the location of mobile device station X, even when using three remote sensors to perform tri-lateration. In these cases, four or more receiving sensors are employed to perform ToA *multi-lateration*.

Figure 2-3 Time of Arrival (ToA)



ToA techniques are capable of resolving location in two-dimensional as well as three-dimensional planes. 3D resolution can be performed by constructing spherical instead of circular models.

A drawback of the ToA approach is the requirement for precise time synchronization of all stations, especially the mobile device (which can be a daunting challenge for some 802.11 client device implementations). Given the high propagation speeds, very small discrepancies in time synchronization can result in very large errors in location accuracy. For example, a time measurement error as small as 100 nanoseconds can result in a localization error of 30 meters. ToA-based positioning solutions are typically challenged in environments where a large amount of multipath, interference, or noise may exist.

The Global Positioning System (GPS) is an example of a well-known ToA system where precision timing is provided by atomic clocks.

Time Difference of Arrival (TDoA)

Time Difference of Arrival (TDoA) techniques use *relative* time measurements at each receiving sensor in place of absolute time measurements. Because of this, TDoA does not require the use of a synchronized time source at the point of transmission (i.e. the mobile device) in order to resolve timestamps and determine location. With TDoA, a transmission with an unknown starting time is received at various receiving sensors, with only the receivers requiring time synchronization.

TDoA implementations are rooted upon a mathematical concept known as *hyperbolic lateration*. In this approach, at least three time-synchronized receiving sensors are required. In [Figure 2-4](#), assume that when station X transmits a message, this message arrives at receiving sensor A with time T_A and at receiving station B with time T_B . The time difference of arrival for this message is calculated between the locations of sensors B and A as the positive constant k , such that:

$$\text{TDoA}_{B-A} = |T_B - T_A| = k$$

The value of TDoA_{B-A} can be used to construct a hyperbola with foci at the locations of both receiving sensors A and B. This hyperbola represents the locus of all the points in the x-y plane, the difference of whose distances from the two foci is equal to $k(c)$ meters. Mathematically, this represents all possible locations of mobile device X such that:

$$|D_{XB} - D_{XA}| = k(c)$$

The probable location of mobile station X can then be represented by a point along this hyperbola. To further resolve the location of station X, a third receiving sensor at location C is used to calculate the message time difference of arrival between sensors C and A, or:


$$\text{TDoA}_{C-A} = |T_C - T_A| = k_1$$

Knowledge of constant k_1 allows for the construction of a second hyperbola representing the locus of all the points in the x-y plane, the difference of whose distances from the two foci (that is, the two receiving sensors A and C) is equal to $k_1(c)$ meters. Mathematically, this can be seen as representing all possible locations of mobile device X such that:

$$|D_{XC} - D_{XA}| = k_1(c)$$

Figure 2-4 illustrates how the intersection of the two hyperbolas $TDoA_{C-A}$ and $TDoA_{B-A}$ is used to resolve the position of station X.

Figure 2-4 Time Difference of Arrival (TDoA)



A fourth receiving sensor and third hyperbola may be added as an enhancement to perform TDoA *hyperbolic multi-ilateration*. This may be required to solve for cases where there may be more than one solution when using TDoA hyperbolic tri-ilateration.

Modern TDoA system designers have derived methods of coping with local clock oscillator drift that are intended to avoid the strict requirement for precision time synchronization of TDoA receivers. For example, time adjustments can be calculated periodically with regard to a reference clock source. These clock adjustments can then be used to correct for offsets from the reference clock elsewhere in the system. In the case of TDoA receivers that are capable of transmitting packets (for example, a TDoA receiver that may be integrated into an 802.11 WLAN access point), another innovative approach may involve the periodic exchange of “timing” packets between receivers. In this approach, time offsets between each receiver and a “reference receiver” can be quantized, with the resulting time adjustment applied accordingly within the system.

Airport ranging systems are a well-known example of TDoA systems in use today. In the world of cellular telephony, TDoA is also referred to as Enhanced Observed Time Difference (E-OTD), and in this specific application offers an outdoor accuracy in that application of about 60 meters in rural areas and 200 meters in RF-heavy urban areas.

ToA and TDoA have several similarities. Both have proven to be highly suitable for large-scale outdoor positioning systems. In addition, good results have been obtained from ToA and TDoA systems in semi-outdoor environments such as amphitheaters and stadiums, as well as contained outdoor environments such as car rental and new car lots or ports of entry. Indoors, TDoA systems exhibit their best performance in buildings that are large and relatively open, with low levels of overall obstruction and high ceilings that afford large areas of clearance between building contents and the interior ceiling. It is precisely in these open, spacious environments that TDoA and ToA-based systems operate at their peak efficiency and performance.

Received Signal Strength (RSS)

Thus far we have discussed two lateration techniques (ToA and TDoA) that use elapsed time to measure distance. Lateration can also be performed by using received signal strength (RSS) in place of time. With this approach, RSS is measured by either the mobile device or the receiving sensor. Knowledge of the transmitter output power, cable losses, and antenna gains as well as the appropriate path loss model allows you to solve for the distance between the two stations.

The following is an example of a common path loss model used for indoor propagation:

$$PL = PL_{1Meter} + 10\log(d^n) + s$$

In this model:

- PL represents the total *path loss* experienced between the receiver and sender in dB. This will typically be a value greater than or equal to zero.
- PL_{1Meter} represents the *reference path loss* in dB for the desired frequency when the receiver-to-transmitter distance is 1 meter. This must be specified as a value greater than or equal to zero.
- d represents the *distance* between the transmitter and receiver in meters.
- n represents the *path loss exponent* for the environment.
- s represents the standard deviation associated with the degree of *shadow fading* present in the environment, in dB. This must be specified as a value greater than or equal to zero.

Path loss (PL) is the difference between the level of the transmitted signal, measured at face of the transmitting antenna, and the level at of the received signal, measured at the face of the receiving antenna. Path loss does not take antenna gains or cable losses into consideration. Path loss represents the level of signal attenuation present in the environment due to the effects of free space propagation, reflection, diffraction, and scattering.

The path loss exponent (n) indicates the rate at which the path loss increases with distance. The value of path loss exponent depends on frequency and environment, and is highly dependent on the degree of obstruction (or “clutter”) present in the environment. Common path loss exponents range from a value of 2 for open free space to values greater than 2 in environments where obstructions are present. A typical path loss exponent for an indoor office environment may be 3.5, a dense commercial or industrial environment 3.7 to 4.0 and a dense home environment might be as high as 4.5.

The standard deviation of shadow fading (s) represents a measure of signal strength variability, (sometimes referred to as “noise”) from sources that are not accounted for in the aforementioned path loss equation. This include factors such as attenuation due to the number of obstructions present, orientation differences between location receiver antennas and the antennas of client devices, reflections due to multipath, and so on. Diversity antenna implementations reduce perceived signal variation due to shadow fading, and for this reason diversity antennas are almost universally recommended. In many indoor installations using diversity antennas, the standard deviation of shadow fading is often seen between 3 and 7 dB.

The generally accepted method to calculate receiver signal strength given known quantities for transmit power, path loss, antenna gain, and cable losses is as follows:

$$RX_{PWR} = TX_{PWR} - LOSS_{TX} + Gain_{TX} - PL + Gain_{RX} - LOSS_{RX}$$

We can directly substitute our equation for path loss into the equation above. This enables us to solve for distance d as follows:

$$d = 10^{\frac{TX_{PWR} - RX_{PWR} - Loss_{TX} + Gain_{TX} - PL_{1meter} + s + Gain_{RX} - Loss_{RX}}{10n}}$$

where the meaning of the terms in the equation above are:

- Rx_{PWR} represents the detected receive signal strength in dB.
- Tx_{PWR} represents the transmitter output power in dB.
- $Loss_{TX}$ represents the sum of all transmit-side cable and connector losses in dB.
- $Gain_{TX}$ represents the transmit-side antenna gain in dBi.
- $Loss_{RX}$ represents the sum of all receive-side cable and connector losses in dB.
- $Gain_{RX}$ represents the receive-side antenna gain in dBi.

Note that all of these are to be specified as positive values.

Solving for distance between the receiver and mobile device allows a circular area to be plotted around the location of the receiver, using the distance d as the radius. The location of the mobile device is believed to be somewhere on this circular plot. As in other techniques, input from other receivers in other cells (in this case, signal strength information or RSSI) can be used to perform RSS *tri-lateration* or RSS *multi-lateration* to further refine location accuracy.

The signal strength information used to determine position can be obtained from one of two sources:

- The network infrastructure reporting the received signal strength at which it receives mobile device transmissions (“network-side”)
- The mobile device reporting the signal strength at which it receives transmissions from the network (“client-side”)

In 802.11 WLANs, the granularity with which RSSI is reported typically varies from radio vendor to radio vendor. In fact, 802.11 client devices produced by different silicon manufacturers may report received signal strength using inconsistent metrics. This can result in degraded and inconsistent location tracking performance. Location tracking solutions that utilize “network-side” RSSI measurements avoid this potential pitfall when supporting mobile devices from various manufacturers, since all measurement of RSSI is performed at the network infrastructure, not at the mobile device. This is a straightforward approach and is approach most often implemented by vendors of RSS lateration solutions, since a much higher degree of control is typically exercised over consistency in network infrastructure versus end user client mobile devices.

Location tracking solutions that rely on “client-side” RSSI measurements must take extra steps to avoid location inaccuracies that may be due to inconsistent mobile device hardware. Since it is not realistic to assume that every mobile device will be provided by the same hardware vendor, a method of “equalizing” any variations in relation to some assumed “reference model” is necessary. For example, assume that a particular positioning system expects to see reported RSSI in a range from -127dBm to +127dBm in 254 increments of 1 dBm each. Mathematical compensation will be required if only some mobile devices in the system can support this expectation (for example, other devices in the system may only be able to report RSSI in a range from -111dBm to +111 dBm in 74 increments of 3dBm each).

Typically, the responsibility for providing such equalization lies with the provider of the location solution. It is common to see such adjustments made through proprietary client software that installed on each mobile device in order to ensure all mobile devices can be located with approximate equal consistency.

To date, implementations using RSS lateration have enjoyed a cost advantage by not requiring specialized hardware at the mobile device or network infrastructure locations. This makes signal strength-based lateration techniques very attractive from a cost-performance standpoint to designers of 802.11-based WLAN systems wishing to offer integrated lateration-based positioning solutions. However, a known drawback to “pure” RSS lateration is that propagation anomalies brought about by anisotropic conditions in the environment may degrade accuracy significantly. This is due in part because in reality, propagation in any cell is far from a purely circular pattern based on an ideal path loss model. “Textbook” theoretical RSS lateration models in their purest form do not provide for the measurement or consideration of variations seen within actual sites, typically assuming only well-known values for path loss and shadow fading.


Pure RSS-based lateration techniques that do not take additional steps to account for attenuation and multipath in the environment rarely produce acceptable results except in very controlled situations. This includes those controlled situations where there is always established clear line-of-sight between the mobile device and the receiving sensors, with little attenuation to be concerned other than free-space path loss and minor impact from multipath.

Angle-Based (Angulation) Techniques

Angle of Arrival (AoA)

The *Angle of Arrival (AoA)* technique, sometimes referred to as *Direction of Arrival (DoA)*, locates the mobile station by determining the *angle of incidence* at which signals arrive at the receiving sensor. Geometric relationships can then be used to estimate location from the intersection of two lines of bearing (LoBs) formed by a radial line to each receiving sensor, as illustrated in [Figure 2-5](#). In a two-dimensional plane, at least two receiving sensors are required for location estimation with improved accuracy coming from at least three or more receiving sensors (*triangulation*).

Figure 2-5 Angle of Arrival (AoA)



In its purest form (that is, where clear line-of-sight is evident between the mobile device X and receiving sensors A and B), mechanically-agile directional antennas deployed at the receiving sensors are adjusted to the point of highest signal strength. The positioning of the directional antennas can be directly used to determine the LoBs and measure the angles of incidence θ_A and θ_B .

In practical commercial and military implementations of AoA, multiple element antenna arrays are used to sample the receiving signal, thereby eliminating the need for more complex and maintenance-intensive mechanical antenna systems. Electronic switching can be performed between arrays or portions of each array, and mathematical computations handled by a background computing system used to extract the angles of incidence. This technique actually involves calculating TDoA between elements of the array by measuring the difference in received phase at each element. In a properly constructed array, there is a small but discernible per element arrival time and a difference in phase. Sometimes referred to as “reverse beam-forming”, this technique involves directly measuring the arrival time of the signal at each element, computing the TDoA between array elements, and converting this information to an AoA measurement. This is made possible because of the fact that in beam-forming, the signal from each element is time-delayed (phase shifted) to “steer” the gain of the antenna array.

A well-known implementation of AoA is the VOR (VHF Omnidirectional Range) system used for aircraft navigation from 108.1 to 117.95 MHz. VOR beacons around the United States and elsewhere transmit multiple VHF “radials” with each radial emanating at a different angle of incidence. The VOR receiver in an aircraft can determine the radial on which the aircraft is situated as it is approaching the VOR beacon and thus its angle of incidence with respect to the beacon. Using a minimum of two VOR beacons, the aircraft navigator is able to use onboard AoA ranging equipment to conduct angulation (or tri-angulation if using three VOR beacons) and accurately determine the position of the aircraft.

AoA techniques have also been applied in the cellular industry in early efforts to provide location tracking services for mobile phone users. This was primarily intended to comply with regulations requiring cell systems to report the location of a user placing an emergency (911) call. Multiple tower sites calculate the AoA of the signal of the cellular user, and use this information to perform

tri-angulation. That information is relayed to switching processors that calculate the user location and convert the AoA data to latitude and longitude coordinates, which in turn is provided to emergency responder dispatch systems.

A common drawback that AoA shares with some of the other techniques mentioned is its susceptibility to multipath interference. As stated earlier, AoA works well in situations with direct line of sight, but suffers from decreased accuracy and precision when confronted with signal reflections from surrounding objects. Unfortunately, in dense urban areas, AoA becomes barely usable because line of sight to two or more base stations is seldom present.

Location Patterning (Pattern Recognition) Techniques

Location patterning refers to a technique that is based on the sampling and recording of radio signal behavior patterns in specific environments. Technically speaking, a location patterning solution does not require specialized hardware in either the mobile device or the receiving sensor (although at least one well-known location patterning-based RTLS requires proprietary RFID tags and software on each client device to enable “client-side” reporting of RSSI to its location positioning server). Location patterning may be implemented totally in software, which can reduce complexity and cost significantly compared to angulation or purely time-based lateration systems.

Location patterning techniques fundamentally assume the following:

- That each potential device location ideally possesses a distinctly unique RF “signature”. The closer to reality this assumption is, the better the performance of the location patterning solution.
- That each floor or subsection possesses unique signal propagation characteristics. Despite all efforts at identical equipment placement, no two floors, buildings, or campuses are truly identical from the perspective of a pattern recognition RTLS solution.

Although most commercially location patterning solutions typically base such signatures on received signal strength (RSSI), pattern recognition can be extended to include ToA, AoA or TDoA-based RF signatures as well. Deployment of patterning-based positioning systems can typically be divided into two phases:

- Calibration phase
- Operation phase

During the operational phase, solutions based on location patterning rely on the ability to “match” the reported RF signature of a tracked device against the database of RF signatures amassed during the calibration phase. Because the database of recorded RF signatures is meant to be compiled during a representative period in the operation of the site, variations such as attenuation from walls and other objects can be directly accounted for during the calibration phase.

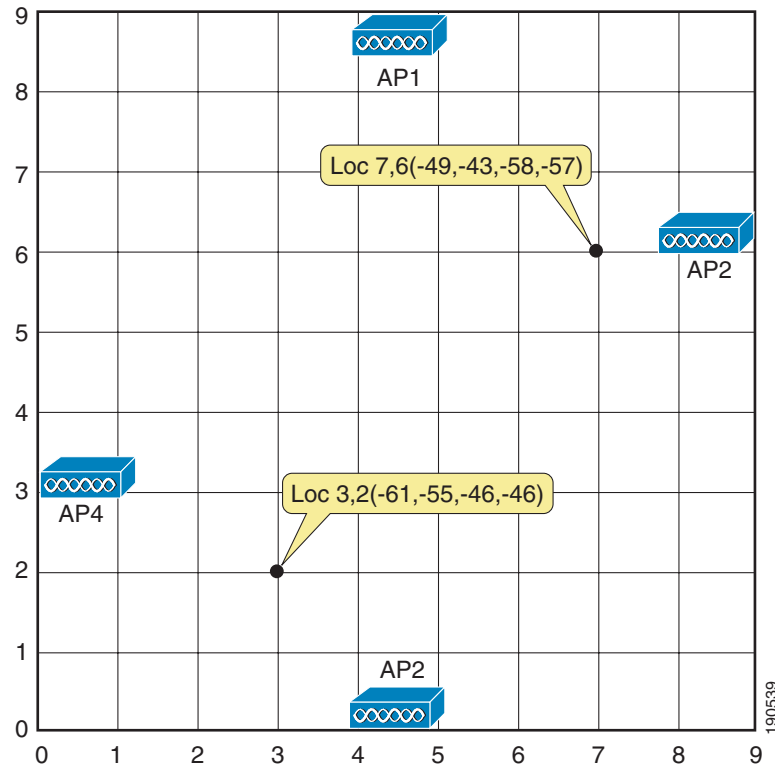
Calibration Phase

During the calibration phase, data is accumulated by performing a walk-around of the target environment with a mobile device and allowing multiple receiving sensors (access points in the case of 802.11 WLANs) to sample the signal strength of the mobile device (this refers to a “network-side” implementation of location patterning).

A graphical representation of the area to be calibrated is typically overlaid with a set of grid points or notations to guide the operator in determining precisely where sample data should be acquired. At each sample location, the array (or *location vector*) of RSS values associated with the calibration device is recorded into a database known as a *radio map* or *training set*. The size of the vector for this sample

location is determined by the number of receiving stations that can detect the mobile device. Figure 2-6 provides a simplified illustration of this approach, showing two sample points and how their respective location vectors might be formed from detected client RSSI.

Figure 2-6 Location Patterning Calibration



Because of fading and other phenomena, the observed signal strength of a mobile device at a particular location is not static but is seen to vary over time. As a result, calibration phase software typically records many samples of signal strength for a mobile device during the actual sampling process. Depending on technique, the actual vector array element recorded may account for this variation via one or more creative approaches. A popular, simple-to-implement method is to represent the array element associated with any specific receiver as the *mean signal strength* of all measurements of that mobile device made by that receiver sensor for the reported sample coordinates. The location vector therefore becomes a vector array of *mean signal strength elements* as shown in the following equation, where x and y represent the reported coordinates of the sample and r represents the reported RSSI:

$$(x, y) = (\bar{r}_{AP1}, \bar{r}_{AP2}, \bar{r}_{AP3}, \bar{r}_{AP4})$$

Operational Phase

In the operational phase, a group of receiving sensors provide signal strength measurements pertaining to a tracked mobile device (network-side reporting implementation) and forwards that information to a location tracking server. The location server uses a complex positioning algorithm and the radio map database to estimate the location of the mobile device. The server then reports the location estimate to the location client application requesting the positioning information.

Location patterning positioning algorithms can be classified into three basic groups:

- *Deterministic algorithms* attempt to find *minimum statistical signal distance* between a detected RSSI location vector and the location vectors of the various calibration sample points. This may or may not be equal to the minimum physical distance between the actual device physical location and the recorded location of the calibration sample. The sample point with the minimum statistical signal distance between itself and the detected location vector is generally regarded as the best raw location estimate contained in the calibration database. Examples of deterministic algorithms are those based on the computation of Euclidean, Manhattan, or Mahalanobis distances.
- *Probabilistic algorithms* use probability inferences to determine the likelihood of a particular location given that a particular location vector array has already been detected. The calibration database itself is considered as an *a priori* conditional probability distribution by the algorithm to determine the likelihood of a particular location occurrence. Examples of such approaches include those using *Bayesian* probability inferences.
- Other techniques go outside the boundaries of deterministic and probabilistic approaches. One such approach involves the assumption that location patterning is far too complex to be analyzed mathematically and requires the application of non-linear discriminant functions for classification (*neural networks*). Another technique, known as *support vector modeling* or *SVM*, is based on risk minimization and combines statistics, machine learning, and the principles of neural networks.

To gain insight into how such location patterning algorithms operate, we can examine a simple example that demonstrates the use of a deterministic algorithm, which in this case will be the Euclidean distance. As stated earlier, deterministic algorithms compute the minimum statistical signal distance, which may or may not be equal to the minimum physical distance between the actual device physical location and the recorded location of the calibration sample.

For example, assume two access points X and Y and a mobile device Z. Access point X reports mobile device Z with an RSS sample of x_j . Almost simultaneously, access point Y reports mobile device Z with an RSS sample of y_j . These two RSS reports can be represented as location vector of (x_j, y_j) . Assume that during the calibration phase, a large population of location vectors of the format $F(x_2, y_2)$ were populated into the location server calibration database, where F represents the actual physical coordinates of the recorded location.

The location server can calculate the Euclidean distance d between the currently reported location vector (x_1, y_1) and each location vector in the calibration radio map as follows:

$$d = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$$

The physical coordinates F associated with the database location vector possessing the minimum Euclidean distance from the reported location vector of the mobile device is generally regarded as being the correct estimate of the position of the mobile device.

In a similar fashion to RSS lateration solutions, real-time location systems using location patterning typically allow vendors to make good use of existing wireless infrastructure. This can often be an advantage over AoA, ToA, and TDoA approaches, depending on the particular implementation. Location patterning solutions are capable of providing very good performance in indoor environments, with a minimum of three reporting receivers required to be in range of mobile devices at all times. Increased accuracy and performance (often well in excess of 5 meters accuracy) is possible when six to ten receivers are in range of the mobile device.

Location patterning applications perform well when there are sufficient array entries per location vector to allow individual locations to be readily distinguishable by the positioning application. However, this requirement can also contribute to some less-than-desirable deployment characteristics. With location

patterning, achieving high performance levels typically requires not only higher numbers of receivers (or access points for 802.11) but also much tighter spacing. In large areas where it is possible for clients to move about almost anywhere, calibration times can be quite long. For this reason, some commercial implementations of location patterning allow the user to segment the target location environment into areas where client movement is likely and those where client movement is possible but significantly less likely, as well as areas where client location is impossible (such as within the thick walls of a tunnel, for example, or suspended within the open air space of an indoor building atrium). The amount of calibration as well as computational resources allocated to these two classes of areas is adjusted by the positioning application according to the relative probability of a client being located there.

The radio maps or calibration databases used by pattern recognition positioning engines tend to be very specific to the areas used in their creation, with little opportunity for re-use. The likelihood is very low that any two areas, no matter how identical they may seem in construction and layout, will yield identical calibration data sets. Because of this, it is not possible to use the same calibration data set for multiple floors of a high-rise office building when using a location patterning solution. This is because despite their similarity, the probability that the location vectors collected at the same positions on each floor being identical is significantly low.

All other variables being equal, location patterning accuracy is typically at its zenith immediately after a calibration. At that time, the information is current and indicative of conditions within the environment. As time progresses and changes occur that affect RF propagation, accuracy can be expected to degrade in accordance with the level of environmental change. For example, in an active logistics shipping and receiving area such as a large scale cross-docking facility, accuracy degradation of 20 percent can reasonably be expected in a thirty day period. Because calibration data maps degrade over time, if a high degree of consistent accuracy is necessary, location patterning solutions require periodic re-verification and possible re-calibration. For example, it is not unreasonable to expect to re-verify calibration data accuracy quarterly and to plan for a complete re-calibration semi-annually.



CHAPTER 3

Cisco Location-Based Services Architecture

This chapter describes the Cisco Location-Based Services (LBS) architecture and has the following main sections:

- [RF Fingerprinting, page 3-1](#)
- [Location-Aware Cisco UWN Architecture, page 3-4](#)
- [Role of the Location Appliance, page 3-7](#)
- [Accuracy and Precision, page 3-11](#)
- [Tracking Clients, Assets and Rogue Devices, page 3-13](#)
- [Cisco Location Control Protocol \(LOCP\), page 3-36](#)

RF Fingerprinting


Cisco *RF Fingerprinting* refers to a new and innovative approach that significantly improves the accuracy and precision of traditional signal strength lateration techniques. Cisco RF Fingerprinting offers the simplicity of an RSSI-based lateration approach with the customized calibration capabilities and indoor performance previously available only in location patterning solutions. RF Fingerprinting significantly enhances RSS lateration by using RF propagation models developed from radio propagation data gathered directly from 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 described for location patterning.

In addition to the use of prepackaged propagation models, RF Fingerprinting offers the ability to develop customized models that are based on on-site data collection. This process allows for the overall attenuation characteristics of the actual environment to be taken into consideration during the derivation of both 2.4 GHz and 5 GHz path loss models. For each calibration grid location, the physical location coordinates of the calibration client (provided by the calibration operator) are recorded along with the client RSSI from three or more LWAPP-enabled access points.

The data accumulated during the calibration phase is statistically processed and groomed, then used to build an RF propagation model used to predict tracked device RSSI around each access point, where the path loss exponent, shadow fading standard deviation, and PL_{1meter} values are calculated from the sample calibration data so as to better reflect specific propagation anomalies present in the environment. This process consists of several computational cycles where the previously-mentioned parameters are calculated for each band. The minimum mean square error (MMSE) estimation technique is used to obtain the *initial* values for the parameters, as shown in [Figure 3-1](#), where the path loss exponent is represented by the slope of the applicable MMSE line of best fit (that is, either default or corrected fit). However, note that in the RF Fingerprinting approach, the selection of a location path loss model does

not end with MMSE. Rather, MMSE is used only as the starting point for the selection of finalized parameters for each band, with the ultimate goal being the optimization of the final path loss model as it pertains to location accuracy. RF Fingerprinting does not rely on good location performance being a by-product of a RF propagation model that simply provides good coverage mapping.

Figure 3-1 MMSE Estimation



To locate a mobile client during the operational phase of RF Fingerprinting, RSS multi-lateration is performed using either a pre-packaged RF model or a customized model created during the calibration phase. This process yields the coordinates of the data point with the highest potential of correctly representing the tracked device's current location. Additional information gleaned from statistical analysis of the distribution of calibration data is then used to further improve location accuracy and precision.

Cisco RF Fingerprinting offers several key advantages over traditional approaches:

- Uses existing LWAPP-enabled Cisco Unified Networking Components—Unlike some other solutions, the location-aware Cisco UWN with RF Fingerprinting provides a Wi-Fi-based RTLS alongside of voice and data services using a combined infrastructure. The Cisco Location Appliance supports location and statistics history and serves as a centralized positioning engine for the simultaneous tracking of up to 2500 devices per appliance. Optional chokepoint triggers can be added to the solution to provide presence and proximity detection if desired, allowing for very granular detection of asset tags, within a range of 25 feet to less than one foot depending on the hardware selected.
- No proprietary client hardware or software required—The location aware Cisco UWN with RF Fingerprinting uses a network-side location model. Because of this, Cisco RF Fingerprinting can provide location tracking for a wide variety of industry-standard Wi-Fi clients (and not just those


with popular Windows-based operating systems) *without the need to load proprietary client-tracking software or location-enabled wireless drivers in each client*. This includes popular VoIP handsets such as the Cisco 7920 and 7921G, devices for which such proprietary add-on location tracking client software is not available.

- Supports popular Wi-Fi active RFID asset tags—Because the location-aware Cisco UWN implements RF Fingerprinting as a network-side model, there is no dependency on proprietary software being resident in RFID asset tags in order to allow for localization. This enables the location-aware Cisco UWN to interoperate with active RFID asset tags from popular vendors including AeroScout, PanGo Networks, WhereNet, G2 Microsystems and others. Asset tags that support the Cisco Compatible Extensions for Wi-Fi Tags specification can take advantage of advanced features introduced with software Release 4.1, such as the ability to pass tag telemetry and chokepoint information to the Cisco UWN. Cisco makes this specification available to Cisco Technology Development Partners (CTDP) and encourages the development of interoperable active RFID tag hardware in compliance with the specification.
- Better accuracy and precision—Cisco RF Fingerprinting yields significantly better performance than solutions employing pure triangulation or RSS 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. With the exception of the calibration phase in location patterning, none of the traditional lateration or angulation approaches discussed thus far take environmental considerations directly into account in this manner. RF Fingerprinting then goes a step further, by applying 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 possibilities and refining accuracy. The net result of these efforts is not only better accuracy but significantly improved precision over traditional solutions.
- Reduced calibration effort—The Cisco RF Fingerprinting technology offers the key advantages of an indoor location patterning solution but with significantly less effort required for system calibration. Although both solutions support on-site calibration, the Cisco RF Fingerprinting approach offers less frequent re-calibration and can operate with a larger inter-access point spacing than location patterning solutions. 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.

Location-Aware Cisco UWN Architecture

The overall architecture of the location-aware Cisco Unified Wireless Network is shown in [Figure 3-2](#).

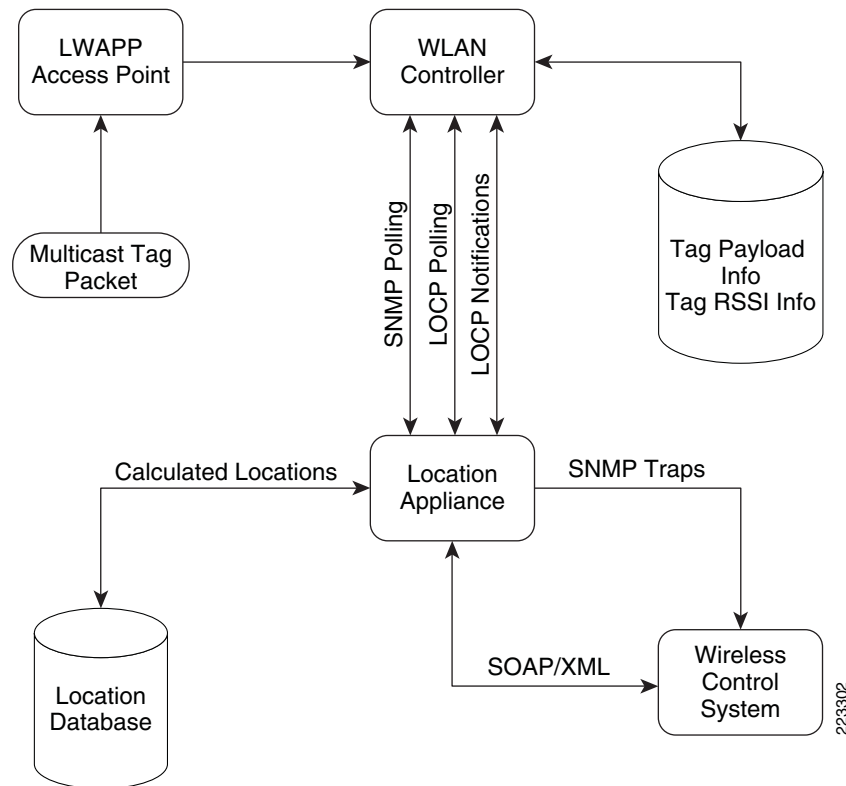
Figure 3-2 Location-Aware Cisco UWN Architecture



Access points forward information to WLAN controllers regarding the detected signal strength of any WLAN clients, asset 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 their regulatory channel set periodically. The collected signal strength information is forwarded to the WLAN controller to which the access point is currently registered, which aggregates the information. The location appliance uses SNMP to poll each controller for the latest signal strength 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 the appropriate controller(s) directly.

A step-by-step flow diagram of this process is provided in [Figure 3-3](#), where the flow of signal strength and tag payload information is shown for active RFID asset tags that communicate via the use of layer two multicasts.

Figure 3-3 Information Flow for Asset Tag RSSI Data



[Figure 3-3](#) summarizes the following events:

-
- Step 1** At each tag transmission interval, the asset tag transmits a multicast frame on each of its configured channels.
- Step 2** At least three access points detect the asset tag's transmission. It is forwarded to the WLAN controller (WLC) to which the detecting access points are registered.
- Step 3** The WLC stores the information payload associated with the asset tag in an internal tag information table indexed by the asset tag MAC address. This information payload can contain information such as battery status and tag or asset telemetry.
- Step 4** For tags detected in the network by access points registered to this WLC, the WLC places the following asset tag information in an internal RSSI table:
- a. Tag MAC address
 - b. AP MAC address
 - c. AP interface
 - d. RSSI measurement
- Step 5** The location appliance periodically polls the WLC for the contents of the tag RSSI table using SNMP.

- Step 6** Commencing with software Release 4.1 of the Cisco UWN, the WLC is polled for the contents of the tag information table using the Cisco Location Control Protocol (LOCP).
- Step 7** The location appliance calculates the location of the asset tag using the RSSI information and stores the location information in its database.
- Step 8** The location appliance dispatches any northbound notifications (such as SNMP traps, emails, syslog or SOAP/XML messages based on the updated asset tag location.
- Step 9** Location end users make use of WCS (or a third party location client) to request location information based on floor maps or search criteria. A request for location information is made from the location to the location appliance via a SOAP/XML online query.
-

Beginning with software Release 4.1 of the location-aware Cisco UWN, LOCP provides for the transmission of asynchronous high-priority messages from the WLAN controller to the location appliance. Included in this category are high-priority tag events such as tag call button alerts, chokepoint proximity and vendor-specific tag payloads.

WCS and the location appliance exchange information (such as calibration maps and network designs) during a process known as *synchronization*. During this process, the partner possessing the more recent information will update the other partner. 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 Cisco Wireless Control System (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, which, as will be further explained in subsequent sections of this document. As a location client, Cisco WCS is capable of displaying a wide multitude of information regarding the current and past location of clients, asset tags, rogue access points, and rogue clients.

**Note**

For important information regarding compatibility between versions of WCS and the Cisco Wireless Location Appliance, refer to *Release Notes for Cisco Wireless Location Appliance Release 3.0* at the following URL:http://www.cisco.com/en/US/products/ps6386/prod_release_notes_list.html.

Location client functionality is not limited to WCS. Third-party applications written in accordance with the Cisco Location Appliance Application Programming Interface (API) can also serve as a location client to the Wireless Location Appliance (as shown in [Figure 3-2](#)). The same information contained in the location appliance that is made available to WCS (including vendor-specific information that may have been received from asset tags) is also made available to third-party same location clients via the location appliance API.

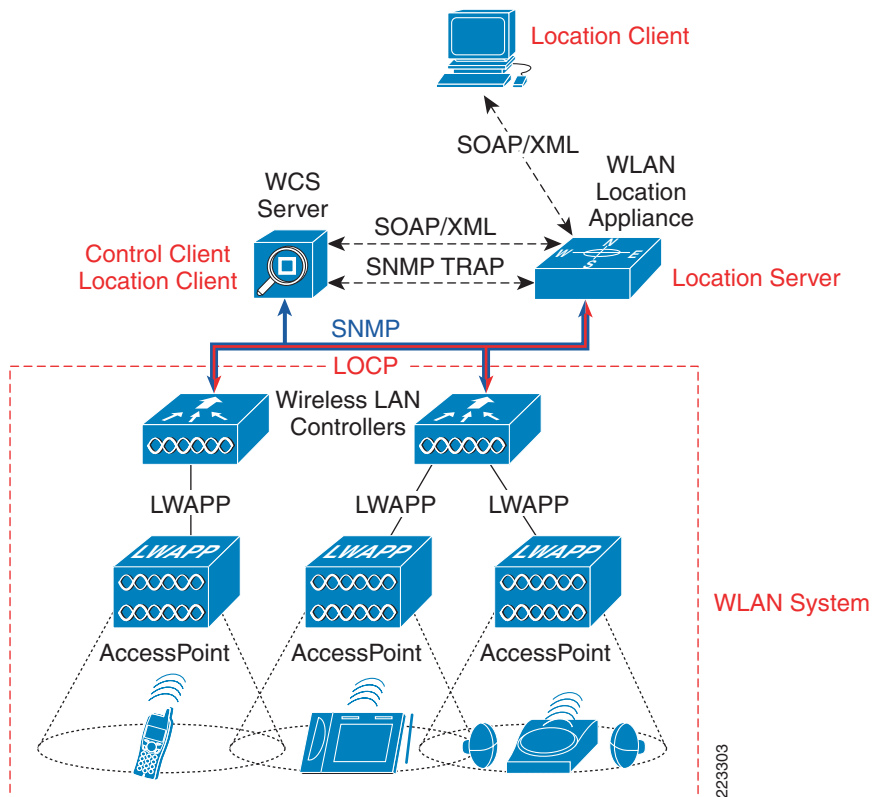
Third-party location clients can synchronize their network designs with the location appliance in a similar fashion to WCS. In this case, the location appliance updates location clients with the latest information regarding network designs and map images. As with WCS, synchronization occurs either on-demand or on a scheduled basis, the timing of which is typically determined by configuration parameters contained within the location client.

The Cisco Location Appliance is also capable of issuing northbound notifications to external systems via email (SMTP), syslog, SNMP traps, or the SOAP/XML protocol. The issuance of these northbound notifications is dependent on the occurrence of one or more of a variety of events, and is discussed in further detail within subsequent sections of this document.

Role of the Location Appliance

The location-aware Cisco UWN can be broken down into four basic component groups, as shown in Figure 3-4.

Figure 3-4 Components of the Location-Aware Cisco UWN



- *Location Client*—The primary role of the location client is to serve as the user interface to the location and asset information contained on the location server. One or more location clients may receive information on a request basis (“pull” mode) or they may assume a listening role awaiting regular transmissions of information from the location server based on pre-defined criteria (“push” mode).
- *WCS*—WCS serves as the default location client to the location appliance, providing location display capabilities that can satisfy most IT-centric and network monitoring requirements. The inherent flexibility afforded by the location appliance API allows for third-party location clients to reside in the UWN in a complementary fashion to WCS. These third-party products may provide a very business-focused UI that concentrates on the management of assets and de-emphasizes the details of RFID, localization and network management.
- *Control Client*—The control client is capable of administering the location server as well as reading or writing all location and configuration data on the location server. In the location-aware Cisco UWN, the role of control client is performed by the Cisco WCS. The control client’s primary function is to populate the server with information about the physical environment (network designs, floors maps, calibration models, access point locations, etc.) and the network elements that should be monitored. The control client may also have management capabilities over one or more of the location servers deployed in the network.

- *Location Server*—The location server provides general location services for a network or part of a network (its *location domain*), and is primarily responsible for running the algorithms that predict client location. The location server may also provide for the storage of historical location information. A location server can communicate with multiple location or control clients. In the location-aware Cisco UWN, the Cisco Wireless Location Appliance fulfills the role of the location server. The Cisco Location Appliance is also capable of issuing notifications to external systems via email (SMTP), syslog, SNMP traps or the SOAP/XML protocol.
- *Wireless LAN System*—The wireless LAN system is comprised of:
 - Embedded software contained within WLAN controllers that functions as an aggregation point for information regarding station/tag/rogue discovery, device tracking and statistics.
 - The mobile devices (tags, mobile stations, rogue clients and rogue access points) that interact with the wireless network and whose location the location-aware Cisco UWN will monitor.
 - Optional infrastructure components, such as chokepoint triggers, that enhance the functionality available from active RFID tags and allow for increased granularity in the localization of these asset tags.

Although it is possible to access the location appliance directly via a console session, all end-user interaction with the location appliance is typically via WCS or a third-party location client application.

The integration of a Cisco Location Appliance into a Cisco Unified Wireless Network architecture immediately enables location improvements over and above the baseline capabilities of the Cisco UWN such as:

- *Scalability*—Adding a Cisco Location Appliance greatly increases the scalability of the location-aware Cisco UWN 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). To handle situations requiring tracking of more than 2500 devices in the enterprise¹, additional location appliances can be deployed. The design can then be partitioned by assigning specific controllers to each appliance. Each appliance is responsible for tracking up to 2500 devices for the controllers and access points within its location domain, and may be managed by a common WCS.
- *Chokepoint Location*—The addition of a Cisco Location Appliance under software Release 4.1 or subsequent releases allows for the use of optional chokepoint triggers from Cisco technology partners such as AeroScout and WhereNet. These devices can assist in providing very granular asset tag location within a range of less than one foot to over twenty feet.
- *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 email (SMTP), syslog, SNMP traps, and SOAP/XML directly to specified destinations. These notifications can be triggered if the client or asset:
 - Changes location.
 - Strays beyond a set distance from pre-determined marker locations.
 - Becomes missing or enters/leaves coverage areas.
 - Experiences a change in battery level.
 - Enters the “stimulation zone” of a chokepoint trigger.
 - Experiences one or more priority conditions, such as:
 - Depression of a tag call button.
 - Detachment of a tag from its asset.

1. If tracked devices roam between location domains, the aggregate tracked device capacity may be reduced.

- An attempt at internal tampering.
- 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 a SOAP/XML interface. Developers can access the Cisco Wireless Location Appliance provisioning services and exchange data in XML format. The location appliance API is available to the Cisco development community along with tools to facilitate solution development. Integration support is available via the Cisco Developer Services Program, a subscription-based service.

**Note**


Complete details on the Cisco Developer Service Program may be found at the Cisco Developer Support website, located at the following URL:

http://www.cisco.com/en/US/products/svcs/ps3034/ps5408/ps5418/serv_home.html.

Location Tracking without a Location Appliance

In order to access any RF Fingerprinting-based location tracking features in the Cisco UWN or even to configure the Cisco Wireless Location Appliance, the Cisco WCS must be appropriately licensed for location usage. When a location-licensed version of WCS is used (verified using WCS main menu bar option Help > About the Software, [Figure 3-5](#)), RF Fingerprinting techniques are used to determine non-chokepoint-based location. When a location appliance is not used with a location-licensed version of WCS, RF Fingerprinting techniques are still used to determine location of tracked devices, but only on-demand and only for a single tracked device at a time.

Figure 3-5 *WCS Licensed for Location*



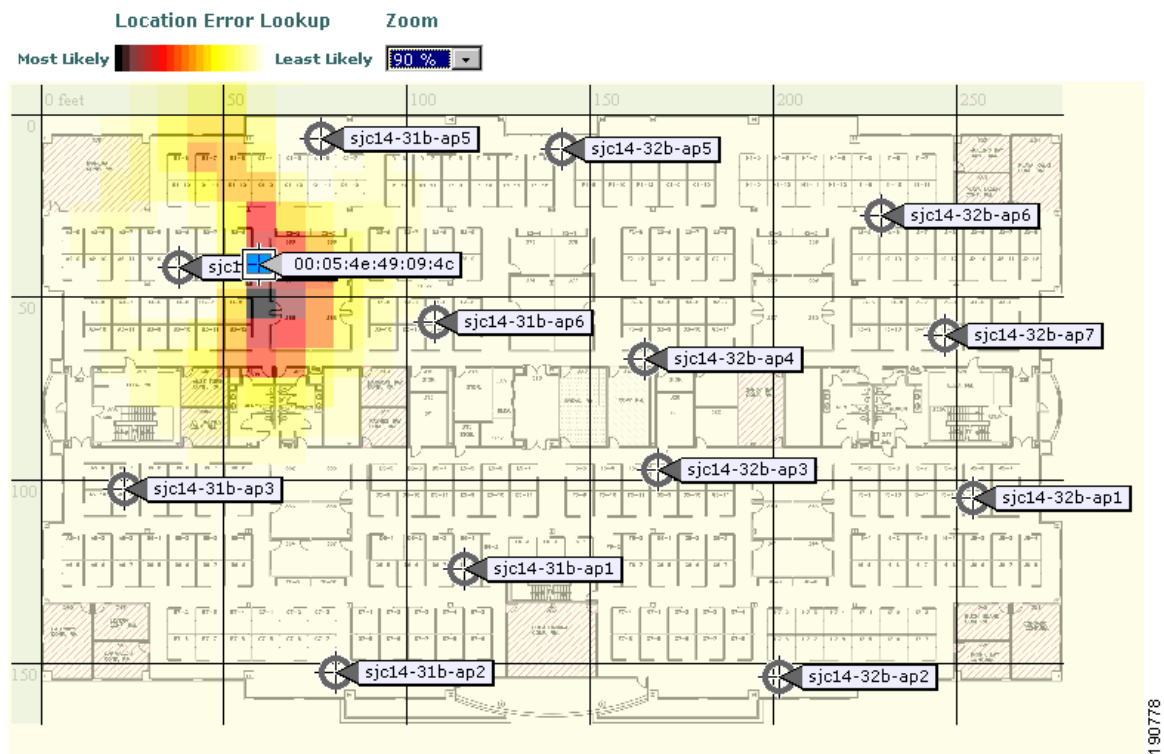
If a location-licensed WCS is used without a location appliance, the following capabilities will be unavailable:

- Ability to configure any Cisco Wireless Location Appliances
- Historical accumulation and playback of location data.
- Tag telemetry and high-priority notifications.
- Chokepoint location.
- The capability to interface to external third-party applications via the SOAP/XML API.
- Simultaneous tracking of multiple devices on a floor map. Location tracking services will be available only as an on-demand service and only for a single device at a time. [Figure 3-6](#) illustrates the use of on-demand localization for a single WLAN client. When using on-demand localization in this manner, it should be noted that colors surrounding the device icon provide an idea of the degree

of location error associated with the icon placement. The darker colors surrounding the icon represent those areas where confidence is high (the probability is higher that the device is physically located where the icon is placed or within this area). The lighter colors represent those areas of lower confidence (the probability is lower that the device is physically located within these areas).

Figure 3-6 On-Demand WLAN Client Localization using WCS with Location License

Maps > Cisco SJ - Site 5 > BLD 14 > 3rd floor



If WCS is licensed for only basic functionality as shown in Figure 3-7, RF Fingerprinting is not employed to determine location. Instead, on-demand location for a single WLAN client or rogue device is performed based on the access point that is detecting the mobile device with the highest signal strength (a derivation of the *nearest access point* concept).

Figure 3-7 WCS Licensed for Only Basic Functions



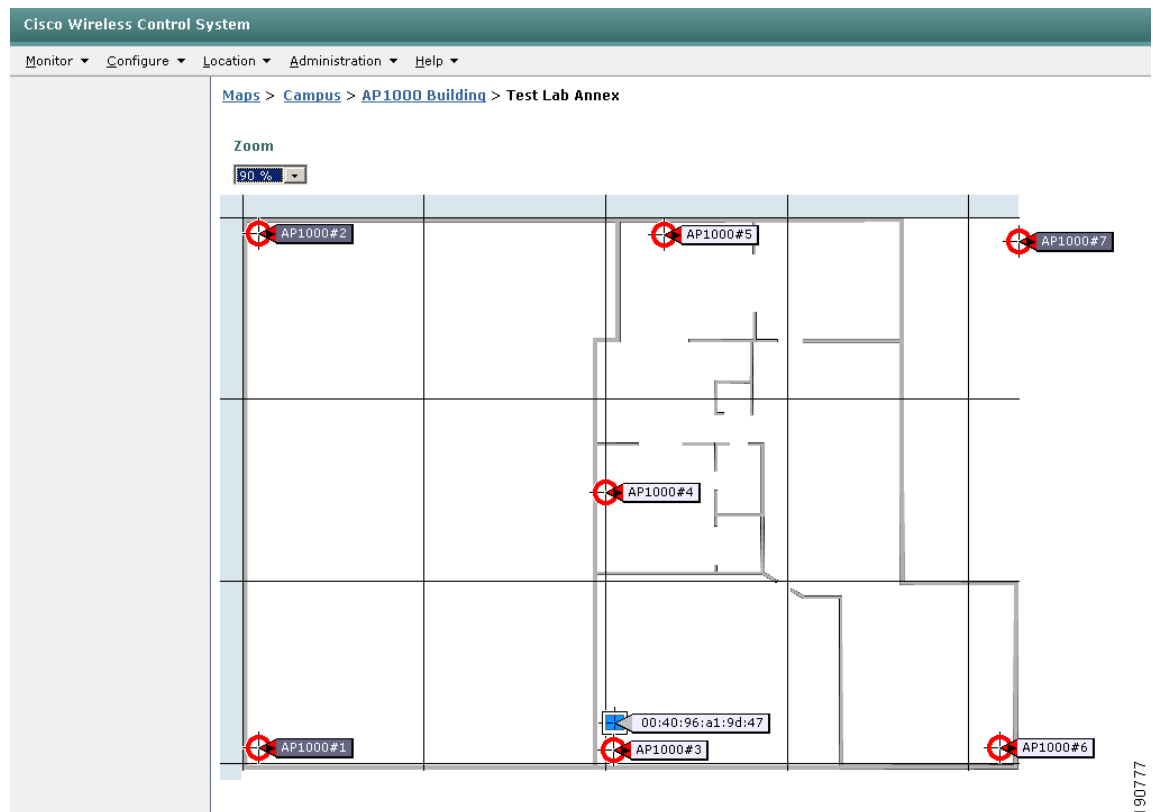
Wireless Control System Version 4.1.91.0

Host: wcswin
Type: Basic
Licensed APs: 50

223305

When using this approach, the tracked device's location is approximated by placing the device icon at the location of the access point detecting it with the highest signal strength, as shown in Figure 3-8. No location probability is displayed in this case.

Figure 3-8 On-Demand Client Localization using WCS with Basic License



Note

A WCS server that is not licensed for location usage cannot be used as a location or control client to the Cisco Wireless Location Appliance. Commencing with software Release 4.1 of the Cisco UWN, on-demand location tracking of asset tags is not possible when using a WCS that is not licensed for location use.

Accuracy and Precision

For most users, the performance metric having the most familiarity and significance is *accuracy*, which typically refers to the quality of the information you are receiving. *Location accuracy* refers specifically to the quantifiable error distance between the estimated and the actual location of a tracked device.

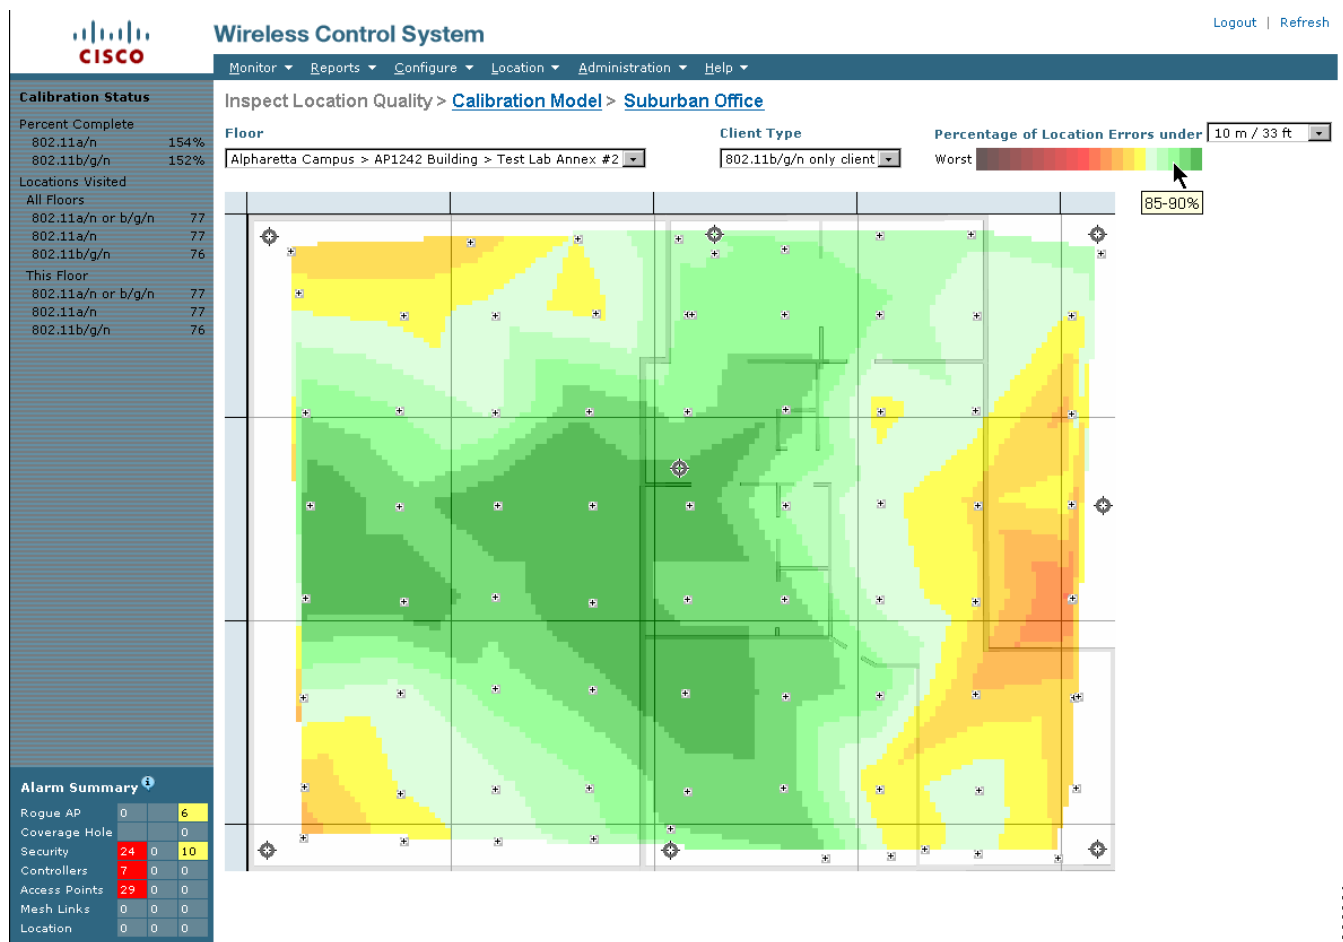
In most real-world applications, however, a statement of location accuracy has little value without the ability of the solution to repeatedly and reliably perform at this level. *Precision* is a direct measure reflecting on the reproducibility of the stated location accuracy. Any indication of location accuracy should therefore include an indication of the confidence interval or percentage of successful location detection as well, otherwise known as the *location precision*.

With deployment in accordance with the best practices outlined in this document, the location-aware Cisco UWN is capable of meeting a baseline performance specification of at least 10 meters accuracy with 90 percent precision. When combined with chokepoint location support, this level of performance can be increased for asset tags possessing chokepoint location capabilities and compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification. Depending on the configured range of the specific chokepoint trigger deployed, a location resolution radius of as little as one foot is possible.

This location appliance's baseline performance level can be achieved by following the best practices along with the use of the design, calibration, and deployment tools described in this and other reference documents. These tools would include predictive, pre-deployment tools such as the *Location Planning* and *Location Readiness* utilities as well as post-deployment tools such as the *Location Inspection* tool.

In order to determine those areas where baseline performance improvements may be necessary, the Location Inspection tool (shown in Figure 3-9), can be used to evaluate what the current, post-calibration levels of accuracy and precision are in the environment. The Location Inspection tool displays (in color coded format) the level of precision at any point from 0 to 5 percent to a maximum of 95 to 100 percent. After viewing the output, the system designer can work with the installation and deployment teams to address any areas requiring remedial attention if necessary.

Figure 3-9 Post-Calibration Location Inspection



223906

Using these tools, the system architect as well as the installation team can not only plan towards the achievement of stated performance goals, but can verify that these targets are indeed being met.

For those interested in a professional service offering that includes the tuning of location performance and much more, Cisco offers a Wireless LAN Location Planning and Design Professional Service. This service offering enlists the skills of trained WLAN engineers to deliver an integrated solution that includes the services Cisco has identified as essential for successful deployment of a secure location-based services solution.

Further information on Cisco Wireless LAN Location Planning and Design Professional Services may be found located at the Cisco Wireless LAN Services website, which is located at the following URL:

http://www.cisco.com/en/US/products/ps8306/serv_home.html

Tracking Clients, Assets and Rogue Devices

This section discusses the mechanics behind the WLAN client probing mechanism and explains how variations in client probing can affect location accuracy. In addition, Cisco Compatible Extensions Location Measurements are explained in detail, along with a close examination of how location for each of the different device categories are displayed by the location client present within WCS.

Client Probing

Fundamentally, the location of WLAN clients is determined based on the RSSI of probe requests detected by access points and forwarded via their registered WLAN controllers to the location appliance. Therefore, the probing behavior of the WLAN and rogue clients in your network can be expected to have a significant impact on the ability of the location appliance to provide accurate location tracking.

Because consistent and regular probing of the network is so important to good WLAN client location fidelity, it is important to understand the mechanics of the process. The process begins with clients issuing probe requests in order to discover the existence of 802.11 networks in their immediate vicinity. An unassociated client may be seen to generate probe requests quite regularly, while clients that are currently associated to a network will typically be seen to issue probe requests less often. Associated clients periodically check their environment for potential access points and networks that they can roam to through a process called *scanning*. In *active scanning*, the client will issue probe requests to solicit probe responses from any access points in its vicinity. From these responses the client forms a list of potential access point roam candidates. Clients may, however, adopt a listen-only approach and simply note the beacons and probe responses they receive from access points around them, without actually soliciting these responses themselves (*passive scanning*). Clients that use passive scanning to determine potential access point roam candidates do not issue probe requests, hence passive scanning in and of itself does little to promote improved location fidelity. It is not unusual to see some clients use a combination of both techniques.

Since the location-aware Cisco UWN uses client probe requests to determine client location, it logically follows that the more consistent the client is in transmitting probe responses, the better the ability of the system will be to provide accurate location tracking of that client. For example, location accuracy can be degraded if a client:

- Refrains from active scanning for long periods
- Does not transmit probe requests across all channels in use
- Does not transmit probe requests for all configured SSIDs
- Transmits probe requests at power levels that deviate abnormally from that expected by the RTLS

IEEE 802.11 standards leave such areas open for interpretation, which does not lead to consistent probing behavior across vendors. This can have both good and not so good connotations from the standpoint of WLAN client location fidelity in the Cisco UWN. While some clients perform active scans and issue probe quite regularly, others may be seen to probe quite minimally.

Cisco Compatible Extensions Location Measurements

The impact of variations in client probing may be greatly reduced by standardizing on clients that are compliant with the Cisco Compatible Extensions for WLAN Devices specification at version 2 or greater. Compatible clients that support the S36 Radio Measurement Requests¹ introduced in Cisco Compatible Extensions for WLAN Devices specification version 2 will perform active scanning and probe all configured SSIDs upon command. Support of this capability enables clients to participate in features such as Cisco Compatible Extensions Location Measurement. When this feature is enabled, registered lightweight access points broadcast Radio Measurement Request frames to their associated clients (via each enabled SSID and radio interface) at a configurable interval from 60 (default) to 32,400 seconds (see [Figure 3-10](#)).

Each Radio Measurement Request contains a beacon request that elicits compatible clients to respond by transmitting probe requests on the channels specified within the Radio Measurement Request. The consistency inherent to this mechanism helps enhance location accuracy for clients so equipped. Note that in software Release 4.1, DFS channels are not included in Radio Measurement Requests.

Using the WCS or controller GUI, Cisco Compatible Extensions Location Measurement can be enabled or disabled per radio interface type (such as 802.11bg or 802.11a) on each WLAN controller. It can also be enabled or disabled globally across controllers using WCS templates. In some cases, more granular control over the Cisco Compatible Extensions Location measurement parameter may be desired, such as when performing testing in specific areas. To support such cases, the WLAN controller CLI allows the Cisco Compatible Extensions Location Measurement feature to be applied to only specific access points if desired. For more information on configuring the Cisco Compatible Extensions Location Measurement using the WLAN controller CLI, refer to the *Cisco WLAN Controller Configuration Guide* at the following URL:

http://www.cisco.com/en/US/products/ps6366/products_configuration_guide_chapter09186a008082d6c5.html#wp1121089.

1. This is one of the features comprising referred to as the *RF Scanning and Reporting* category of Cisco Compatible Extensions for WLAN Devices. A complete list of Cisco Compatible Extensions features are found at the following URL:

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

Figure 3-10 Enabling CCX Location Measurement Using WCS Controller Template

10.1.56.18 > 802.11b/g Parameters

General		Data Rates	
802.11b/g Network Status	<input checked="" type="checkbox"/> Enabled	1 Mbps	Mandatory
802.11g Support	<input checked="" type="checkbox"/> Enabled	2 Mbps	Mandatory
Beacon Period	100	5.5 Mbps	Mandatory
DTIM Period (beacon intervals)	1	6 Mbps	Supported
Fragmentation Threshold (bytes)	2346	9 Mbps	Supported
Short Preamble *	<input checked="" type="checkbox"/> Enabled	11 Mbps	Mandatory
Pico Cell Mode	<input type="checkbox"/> Enable	12 Mbps	Supported
Template Applied	802.11bConfig_166	18 Mbps	Supported
		24 Mbps	Supported
		36 Mbps	Supported
		48 Mbps	Supported
		54 Mbps	Supported

802.11b/g Power Status		Noise/Interference/Rogue Monitoring Channels	
Dynamic Assignment	Automatic	Channel List	DCA Channels
Current Tx Level	5	CCX Location Measurement	
Control Interval sec	600	Mode	<input checked="" type="checkbox"/> Enabled
Dynamic Tx Power Control	<input checked="" type="checkbox"/> Enabled	Interval (seconds)	60 **

802.11b/g Channel Status	
Assignment Mode	Automatic
Update Interval sec	600
Avoid Foreign AP Interference	<input checked="" type="checkbox"/> Enabled
Avoid Cisco AP load	<input type="checkbox"/> Enabled
Avoid non 802.11 Noise	<input checked="" type="checkbox"/> Enabled
Signal Strength Contribution	<input checked="" type="checkbox"/> Enabled

* Controller must be rebooted for new value to take an effect

Save **Audit**

190542

If the Cisco Compatible Extensions Location Measurement parameter is enabled, Radio Measurement Requests will be broadcast to all associated WLAN clients, regardless of their capability to support Cisco Compatible Extensions. Clients that do not support S36 Radio Measurement Requests (such as those supporting Cisco Compatible Extensions version 1 or those not compatible with the Cisco Compatible Extensions for WLAN Devices specification at all) will ignore any Radio Measurement Requests that are received.

**Note**

When using clients equipped with the Intel® PRO/Wireless 3945ABG Network Connection or the Intel® PRO/Wireless 2915ABG Network Connection adapter, it is important to note that the default “Personal Security” settings of the Intel ProSet Configuration Utility do not include compatibility with the Cisco Compatible Extensions specification. When using this default “personal” level of wireless security (which is not intended for enterprise use), clients equipped with the Intel 3945ABG or 2915ABG client adapters will not support S36 broadcast radio measurement requests and are not compliant with the Cisco Compatible Extensions specification for WLAN devices. In order to enable compatibility with the Cisco Compatible Extensions specification and the support of S36 radio

measurement requests, the Intel ProSet client supplicant must be used to reconfigure the client for “Enterprise Security” and enable Cisco Compatible Extensions. Figure 5-42 on page 5-60 and Figure 5-43 on page 5-61 illustrate how this is performed.

An example of a Radio Measurement Request can be seen in Figure 3-11. This request is seen to emanate from an access point with a 802.11b/g interface MAC address of 00:14:1B:59:42:72.

Figure 3-11 Broadcast Radio Measurement Request

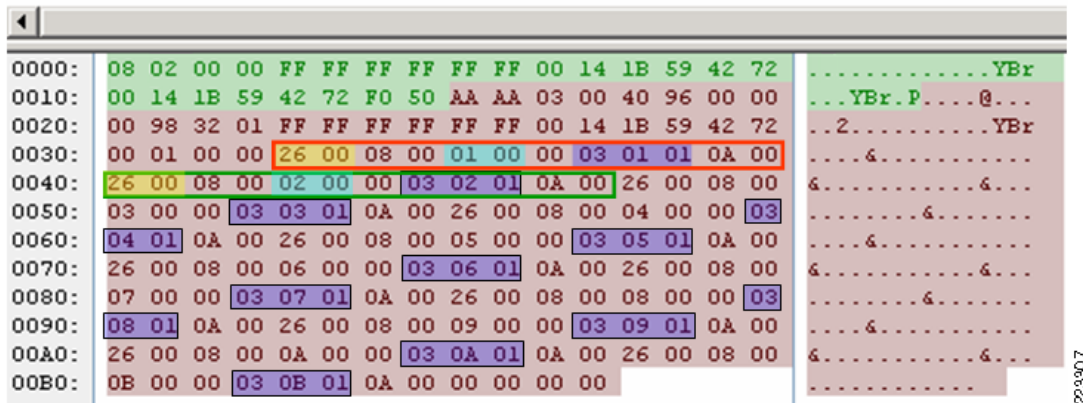
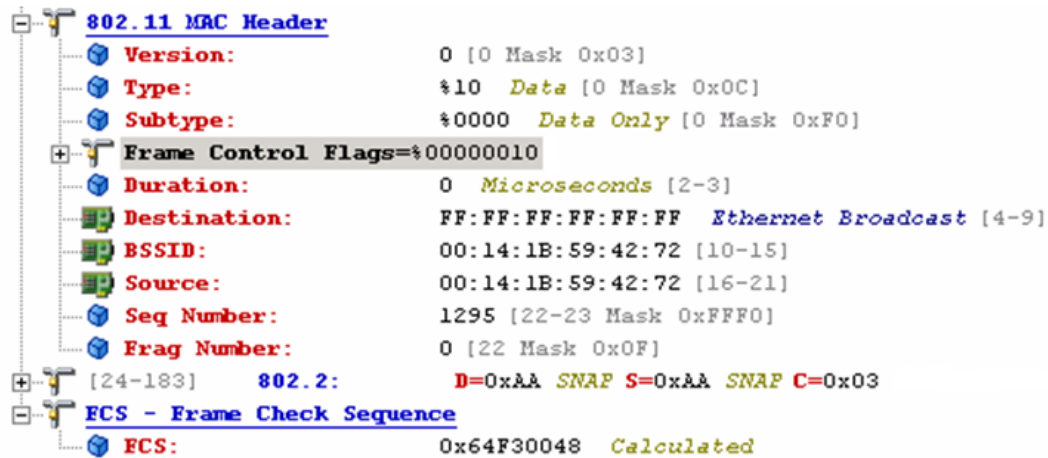


Figure 3-11 provides several key pieces of information that supports our understanding of the effect of Radio Measurement Requests on WLAN clients. We see in Figure 3-11 that the frame broadcast to the associated clients actually contains multiple Radio Measurement Request Elements, the first of which is highlighted within the red rectangle beginning at hex offset 0x0034. Looking closer into the Radio Measure Request Element we see the following¹:

- The element ID of 0x2600 appears at hex offset 0x0034, identifying that what follows is a Measurement Request Element (shown in yellow).
- The measurement token of 0x01 appears at offset 0x0038. This is a non-zero hex value that is unique amongst the Measurement Request Elements in a particular Measurement Request frame.

1. Fields in the radio information elements follow the 802.11 convention of sending the least significant byte first.

- At HEX offset 0x003B, we see the first of several Radio Measurement Request Element detail fields (the first three bytes of each shown highlighted within blue rectangles). Upon closer examination, we can see that each detail field contains:
 - The Measurement Type Definition of 0x03, indicating that this is a Beacon Request. This measurement type requests that the receiving client station perform active or passive scanning (see the Scan Mode Definition field below), and forward the results of those scans upstream to the UWN when it has completed. When performing an active scan, the client device transmits probe request frames that solicit probe responses from receiving access points. In contrast, a passive mode scan requests that 802.11 client devices simply listen for beacon or probe response frames but does not require active solicitation.
 - The Channel Number that the Beacon Request should apply to. This is the second octet of the measurement request detail field and can be seen at hex offset 0x003C with a value of 0x01.
 - The Scan Mode Definition of 0x01, the third octet of the management request detail field. This can be seen at hex offset 0x003D. A Scan Mode Definition of 0x01 indicates that an active scan should be performed. From a strict location fidelity perspective, a passive scan would do little to enhance client location fidelity (since no probe requests are generated). Therefore, when CCX Location Measurement is enabled on the controller, the Scan Mode Definition will always be set to request that active scanning be performed.

Taking all three Radio Measurement Detail fields into consideration, we see that this Radio Measurement Request Element contains a Beacon Request for an active scan to be performed on Channel one.

Note that there are eleven Measurement Request Element fields contained in the Radio Measurement Request. [Figure 3-11](#) highlights only two of them, the first contained within a red rectangle and the second within a green rectangle. This is understandable given that this Radio Measurement Request is being issued on a 802.11b/g radio interface that is operating in the North American regulatory domain with eleven available channels. In the subsequent measurement requests (indicated by the green rectangle), the Channel Number field (seen at hex offset 0x44) is sequentially incremented by 1 from that of the initial measurement request. This continues in the Measurement Request Element fields that follow until the value of 0xB (11) is reached.

802.11bg clients compliant with the Cisco Compatible Extensions for WLAN Devices specification version 2 or greater and supporting S36 Radio Measurement Requests receive the frame shown in [Figure 3-11](#) and will perform an active scan of the specified channels as part of the radio measurement process. When the probe requests are received by access points in the vicinity of such clients, they forward (via their registered controller) signal strength measurements to the location appliance that is used to localize the client. In addition, clients also collect the RSSI information of all probe responses received during the measurement duration, and forward this to the Cisco UWN in a Radio Measurement Report frame.

As per the Cisco Compatible Extensions for WLAN Devices specification version 2, WLAN clients supporting S36 Radio Measurement Requests should:

- Perform radio measurements on the channel over which the Measurement Request was received without significantly degrading performance.
- Perform measurements on non-serving channels while temporarily buffering outgoing traffic.
- Respond to each Radio Measurement Request frame accepted with a Radio Measurement Report frame.
- Disregard any measurement requests that would significantly degrade performance of the client device.
- Support active scanning.

WLAN Clients

Wireless LAN clients and properly configured work-group bridges are displayed on the WCS location floor maps using a blue rectangle icon, as shown in [Figure 3-12](#). To display WLAN clients on the WCS location floor map, ensure that the **Clients** checkbox option is enabled from the **Layers** dropdown selector at the top of the floor map display, and click **Load** in the left-hand column. To avoid excessive clutter, WCS will display the first 250 WLAN clients on the floor map. To view the location of WLAN clients beyond the first 250, client filtering must be used.

Figure 3-12 WCS WLAN Client Location Map



Note that the graphical location information shown can be filtered by WCS based on the age of the information. Thus in [Figure 3-12](#), WCS displays device location information that has aged up to 15 minutes. This value can be set to 2 or 5 minutes if you would like to view location information received more recently, or ½, 1, 3, 6, 12, or 24 hours for information that is older.

By clicking on the blue chevron **>** that is displayed to the right of the **Clients** checkbox option, client filtering options can be specified and additional information retrieved, such as:

- The total number of WLAN clients detected on this floor.
- Small icons (shown in [Figure 3-12](#)) or standard size icons can be selected. When using small icons, descriptive text is not displayed on the floor map for the client except when a mouse-over is performed. When using standard size icons, an on-screen tag is displayed that is configurable for IP address, user name, MAC address, asset name, asset group, or asset category.

- Either all WLAN clients can be displayed, or filtering can be performed to select which clients to display on the floor map. This can be based on IP address, user name, MAC address, asset name, asset group, asset category, or controller. Additional filtering can be specified for SSID and RF protocol (802.11a or 802.11b/g). As mentioned previously, only up to 250 WLAN clients will be shown at on the floor maps at any one time. If there are greater than 250 WLAN clients detected, the total number found will be indicated in the left hand column status area during each communication cycle between WCS and the location appliance. It is recommended that filtering be used to reduce the total number of WLAN clients selected for display if you receive this warning.

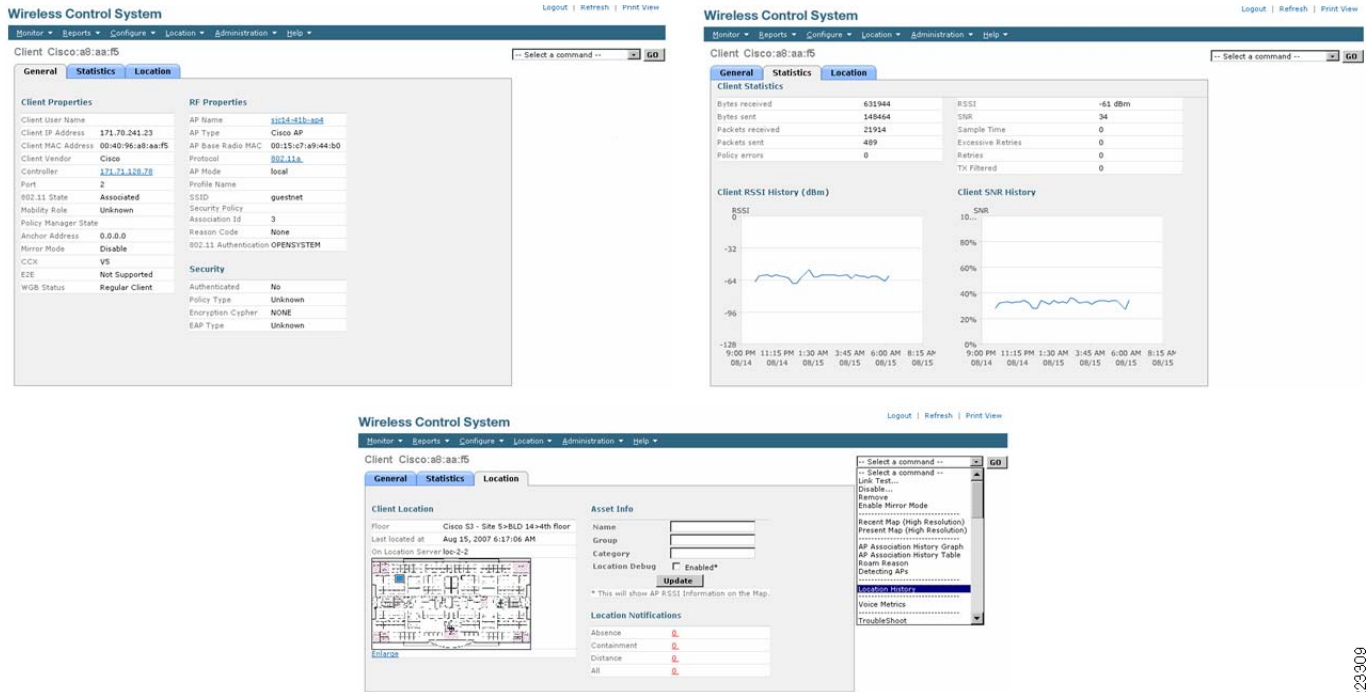
In software Release 4.1 of the Cisco UWN, WLAN controllers provide support for the maximum number of WLAN clients listed in [Table 3-1](#).

Table 3-1 Maximum WLC Client Capacity

Controller Model	WLAN Clients Supported
2006	256
2106	256
4402	2,500
4404	5,000
WiSM	10,000
NM-WLC6	256
NME-WLC8/12	350
3750G	2,500

Complete information on any displayed WLAN client can be obtained simply by left-clicking on the appropriate blue rectangular icon on the floor map, as shown in [Figure 3-13](#). Note that name, group, and category information can be assigned to the client under the “location” submenu, which can then be used to identify the asset on the floor map display.

Figure 3-13 WLAN Client Detailed Information



Note that Figure 3-13 also includes a hyperlinked listing of location notifications as well as a miniature location map showing the client location. By enlarging the map and enabling the Location Debug parameter, WCS displays the last detected RSSI levels of each access point detecting the WLAN client, as shown in Figure 3-14.



Note

The setting of the Location Debug Enable checkbox does not survive a restart of the *locserverd* application or a reboot of the location appliance.

This RSSI information is collected in a similar fashion to that shown by the **show client detail <mac address>** command, and provides an alternative to the CLI command for determining the detected RSSI of WLAN clients (see Figure 3-14).

Figure 3-14 WLAN Client Detected RSSI with Location Debug Enabled

The screenshot displays the Cisco Wireless Control System interface for a client. The client's MAC address is Intel:09:75:aa. The interface is divided into several sections:

- Client Location:** Shows the client's current location as "Alpharetta Campus > AP1242 Building > Test Lab Annex #2". It also indicates the last location time as "Aug 15, 2007 11:36:13 AM" and the location server as "AeS_Loc1". A floor plan map shows the client's location with a blue square. An "Enlarge" link is provided below the map.
- Asset Info:** Includes fields for Name, Group, and Category. The "Location Debug" checkbox is checked and labeled "Enabled*". An "Update" button is present. A note states: "* This will show AP RSSI Information on the Map."
- Location Notifications:** A table showing notification counts:

Absence	0
Containment	0
Distance	0
All	0
- Map View:** A detailed floor plan map showing RSSI readings from various access points. A tooltip for AP 00:14:1b:59:43:80 (AP1242#5) displays the following RSSI Readings table:

Detected RSSI	Radio Type	Age when Located
-53 dBm	11b/g/n	63 secs

Wireless client device location history may be displayed by selecting **Location History** from the dropdown menu at the top right-hand corner of the screen (illustrated in the location screen view of Figure 3-14) and clicking **Go**. Past location history stored within the location appliance is displayed for the wireless client via the screen shown in Figure 3-15.

Figure 3-15 WLAN Client Location History

Wireless Control System

Monitor Reports Configure Location Administration Help

Client Cisco:a8:aa:f5

Client User Name _____ Client MAC Address 00:40:96:a8:aa:f5
 Client IP Address 0.0.0.0 Client Vendor Cisco

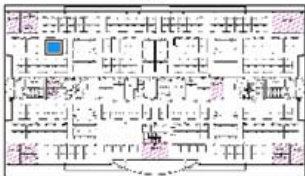
From : Mon Jul 16 03:03:28 EDT 2007
 To : Wed Aug 15 10:35:05 EDT 2007

	Time Stamp	Floor
1	Wed Aug 15 10:35:05 EDT 2007	Cisco SJ - Site 5>BLD 14>4th floor
2	Wed Aug 15 08:35:05 EDT 2007	Cisco SJ - Site 5>BLD 14>4th floor
3	Wed Aug 15 06:35:05 EDT 2007	Cisco SJ - Site 5>BLD 14>4th floor
4	Wed Aug 15 04:35:05 EDT 2007	Cisco SJ - Site 5>BLD 14>4th floor
5	Wed Aug 15 02:35:05 EDT 2007	Cisco SJ - Site 5>BLD 14>4th floor

Change selection every

Client Location

Location Calculated Wed Aug 15 10:26:33 EDT 2007
 Floor Cisco SJ - Site 5>BLD 14>4th floor



[Enlarge](#)

Client Statistics

Data Collected Wed Aug 15 10:21:44 EDT 2007

Bytes received 698622
 Bytes sent 170998
 Packets received 24165
 Packets sent 548
 Policy errors 0
 RSSI -60dBm
 SNR 32

Client Properties

Data Collected Wed Aug 15 10:26:33 EDT 2007

Controller 171.71.128.78
 Port 2
 802.11 State Associated
 Mobility Role Unknown
 Policy Manager State _____
 Anchor Address 0.0.0.0
 CCX V5
 E2E Not Supported

RF Properties

AP Name sjc14-41b-ap5
 AP Type Cisco AP
 AP Base Radio MAC 00:15:c7:a9:43:10
 Protocol 802.11a
 AP Mode local
 SSID _____
 Association Id 10
 Reason Code 0
 802.11 Authentication 0
 Status Code 0
 CF Pollable Not Implemented
 CF Poll Request Not Implemented
 Short Preamble Not Implemented
 PBCC Not Implemented
 Channel Agility Not Implemented
 Timeout 0
 WEP State ENABLE

Security

Authenticated No
 Policy Type Unknown
 Encryption Cypher 5
 EAP Type Unknown

22:33:11

In many cases, it is desirable to sequentially display the location history of a client device in order 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. Cisco WCS and the location appliance make it possible to view each location history record 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, simply click on the **Play** button shown in Figure 3-15.

802.11 Active RFID Tags

The location-aware Cisco UWN readily detects 802.11 Wi-Fi active RFID tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification (such as those from AeroScout, WhereNet, G2 Microsystems and InnerWireless (PanGo), amongst others) and displays them on WCS floor maps using a yellow tag icon, as shown in Figure 3-16. These asset tags typically do not associate to the WLAN infrastructure and are not typically located on the basis of probe requests). Instead, these asset tags transmit messages to the location-aware UWN on a periodic basis using layer two multicasts. If an asset tag has an optional mode that allows for full WLAN association, those tags will be represented on WCS location floor maps as blue rectangles (WLAN clients) during the time they are operating in this mode.

To display the location of asset tags on the WCS location floor map, ensure that the **Clients** checkbox option is enabled from the **Layers** drop down selector at the top of the floor map display, and click **Load** in the left-hand column. To avoid excessive clutter, WCS will display the first 250 asset tags on the floor map. To view the location of asset tags beyond the first 250, asset tag filtering must be used. It is assumed that all other components of the location-aware Cisco UWN have been properly configured to collect asset tag information.

Figure 3-16 RFID Tag Location Map



The graphical location information shown can be filtered by WCS based on the age of the information. In [Figure 3-16](#) WCS displays location appliance information that has aged up to 15 minutes. This value can be set to 2 or 5 minutes if it is desired to view only very recent location information, or ½, 1, 3, 6, 12, or 24 hours to include information that is older.

By clicking on the blue chevron **>** that is displayed to the right of the **802.11 Tags** checkbox option, tag filtering options and additional information can be displayed, such as:

- The total number of asset tags detected on this floor can be displayed.
- Small icons (shown in [Figure 3-16](#)) or standard size icons can be selected. When using small icons, text is not displayed on the floor map for the asset tag except when a mouse-over is performed. When using standard size icons, an on-screen tag is displayed, which is configurable for MAC address, asset name, asset group, or asset category.
- Either all asset tags can be displayed or filtering can be performed to select which asset tags to display on the floor map. This can be based on MAC address, asset name, asset group, asset category, or controller. As mentioned previously, only up to 250 asset tags will be shown on the floor maps at any one time. If there are greater than 250 asset tags detected, the total number found will be indicated in the left hand column status area during each communication cycle between WCS and the location appliance. It is recommended that filtering be used to reduce the total number of asset tags selected for display if you receive this warning.

In software Release 4.1 of the Cisco UWN, WLAN controllers provide support for the maximum number of asset tags listed in [Table 3-2](#).

Table 3-2 Maximum WLC Asset Tag Capacity

Controller Model	Asset Tags Supported
2006	500
2106	500
4402	1250
4404	2500
WiSM	5000
NM-WLC6	500
NME-WLC8/12	500
3750G	1250

Complete information on any displayed asset tag can be obtained by clicking on the yellow tag icon associated with the tag. WCS responds with the information shown in [Figure 3-17](#). Beginning with software Release 4.1 of the location-aware Cisco UWN, tag telemetry, chokepoint and tag status information are also displayed on the Tag Details screen shown in [Figure 3-17](#), along with enhanced battery reporting information.

Figure 3-17 RFID Tag Detailed Information

Wireless Control System

[Tags](#) > Aeroscourt Tag 00:0c:cc:5c:05:17

-- Select a command -- **GO**
 -- Select a command --
 Location History

Tag Properties

Vendor	Aeroscourt
Controller	10.1.96.18
Battery Life	Batt remaining = 80 %, Days remaining = 0, Tolerance = +/- 20 %, Battery Age = 0

Location

Floor	Alpharetta Campus>AP1242 Building>Test Lab Annex #2
Last located at	Aug 15, 2007 4:16:35 PM
On Location Server	AeS_Loc1
Last Chokepoint	00:0c:cc:60:1e:8a
Chokepoint Encountered	Wed Aug 15 16:00:41 EDT 2007

Asset Info

Name	<input type="text"/>
Group	<input type="text"/>
Category	<input type="text"/>
Location Debug	<input checked="" type="checkbox"/> Enabled*

Update

* This will show AP RSSI Information on the Map.

Statistics

Bytes received	104877
Packets received	2012

Location Notifications

Absence	0
Containment	0
Distance	0
All	0

Telemetry Data

TEMPERATURE	: 100.0 degrees Celsius
QUANTITY	: 29
MOTION	: 29.0576 meters/sec
HUMIDITY	: 80 %
MOTIONPROB	: Acceleration
FUEL	: 29.0576 liters

Emergency Data


Reason:	Panic Button Pressed
---------	----------------------

[Enlarge](#)

223313

In some cases the location appliance may place two or more asset tags at the same predicted location, such that any attempt to graphically represent them as individual icons would result in almost complete overlap. A tag summary icon (a yellow tag with black horizontal lines) is used to resolve such situations, as shown in [Figure 3-18](#).

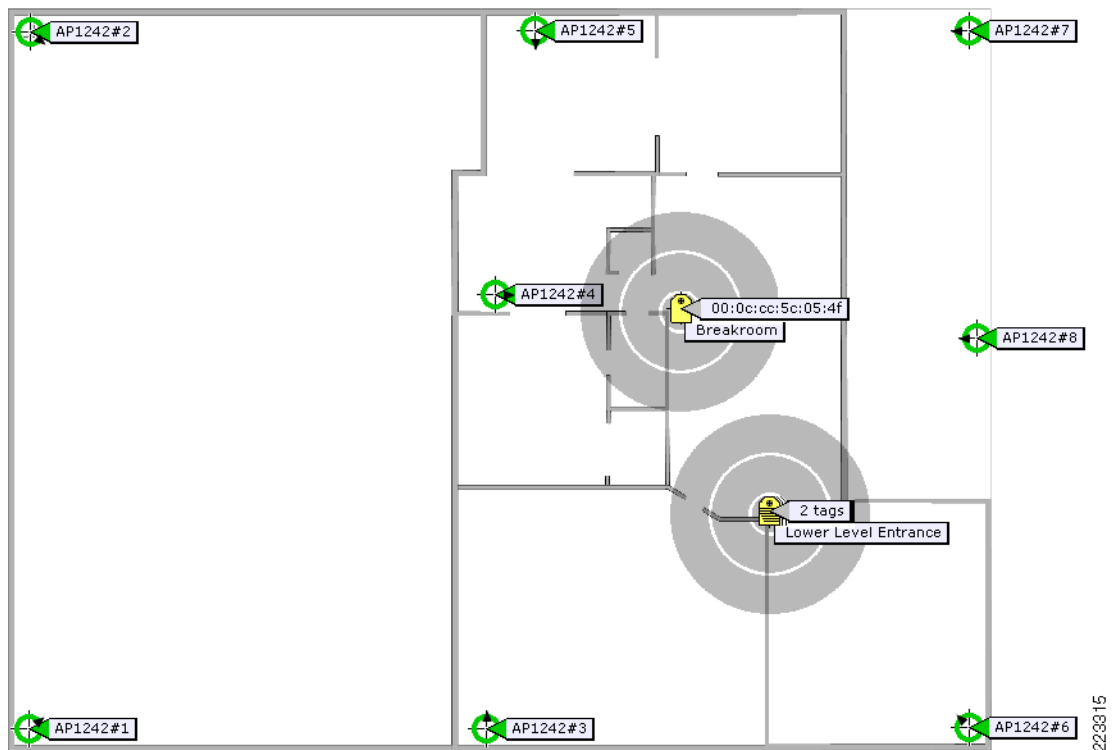
Figure 3-18 Tag Summary Icon and Summary Descriptor



Performing a mouse-over of the tag summary icon brings up a tag summary descriptor shown, which summarizes pertinent tag characteristics. Clicking on “Next” scrolls through the descriptor information for each tag MAC address at this location, and clicking on “Details” at any time brings up the Tag Details panel shown in [Figure 3-17](#).


The tag summary icon becomes especially useful when *chokepoint location* (introduced with software Release 4.1 of the Cisco UWN) is used. When chokepoints have been defined to the system and properly defined on floor maps, the icon of any asset tag that known to be within range of the chokepoint trigger will be placed at the center of the chokepoint icon (shown in [Figure 3-19](#) at the “Breakroom” chokepoint). However, if more than one asset tag is in proximity of the same chokepoint, the tag icons will overlap and usability will suffer. In this situation, the tag summary icon shown in [Figure 3-18](#) once again is used to restore clarity. An example of the tag summary icon being used in this can be seen in [Figure 3-19](#), at the chokepoint labeled “Lower Level Entrance”.

Figure 3-19 Tag Summary Icon and Chokepoint Location



Note that [Figure 3-20](#) also includes a hyperlinked listing of location notifications as well as a miniature location map of the asset tag's location. By enabling the Location Debug parameter and enlarging the map, WCS displays the last detected RSSI levels of each access point detecting the asset tag. This RSSI information is collected in a similar fashion to that shown by the `show rfid detail <mac address>` command, and provides an alternative to the CLI command for determining the detected RSSI of asset tags. As can be seen in [Figure 3-20](#), additional information regarding the radio type and age of the last detected signal strength reading is available by performing a mouse-over of any access point.

Figure 3-20 Asset Tag Detected RSSI with Location Debug Enabled



Asset tag location history may be displayed by selecting **Location History** from the dropdown menu at the top right-hand corner of the screen shown in Figure 3-20 and then clicking on **Go**. Past location history stored within the location appliance will be displayed for the asset tag, along with last values recorded for location statistics, tag telemetry, battery and “emergency” status, as shown in Figure 3-21.

Figure 3-21 Asset Tag Location History

Wireless Control System

Monitor ▾
Reports ▾
Configure ▾
Location ▾
Administration ▾
Help ▾

Aer scout Tag 00:0c:cc:5c:05:17 -- Select a command -- ▾ **GO**

Asset Name	Asset Group
Asset Category	MAC Address 00:0c:cc:5c:05:17

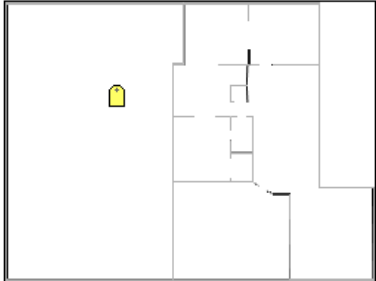
From : Wed Aug 15 17:10:37 EDT 2007
To : Wed Aug 15 19:07:40 EDT 2007

	Time Stamp	Floor	Battery Status
1	Wed Aug 15 19:07:40 EDT 2007	Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %
2	Wed Aug 15 19:06:40 EDT 2007	Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %
3	Wed Aug 15 19:05:40 EDT 2007	Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %

Change selection every 2 secs ▾ Play Stop

Location

Location Calculated	Wed Aug 15 19:08:39 EDT 2007
Floor	Alpharetta Campus>AP1242 Building>Test Lab Annex #2



[Enlarge](#)

Tag Statistics

Data Collected	Wed Aug 15 19:08:36 EDT 2007
Bytes received	162597
Packets received	3122

Telemetry Data

TEMPERATURE	: 100.0 degrees Celsius
QUANTITY	: 29
MOTION	: 29.0576 meters/sec
MOTIONPROB	: No Movement
FUEL	: 29.0576 liters

Emergency Data

Reason:	Panic Button Pressed
Tamper State:	Inactive

Tag Properties

Data Collected	Wed Aug 15 19:07:39 EDT 2007
Controller	10.1.96.18
Battery Status	80 %

223317

In many cases, it is desirable to sequentially display the location history of an asset tag so as to better visualize and trace the movement of the asset tag (and the attached asset) throughout the environment over time. This can be very useful, for example, in establishing a trail of motion in security and monitoring applications. Cisco WCS and the location appliance make it possible to do this by playing back each location history record 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, simply click the **Play** button shown in [Figure 3-21](#) and past location history should start being displayed both in tabular form and graphically. Large amounts of location history data may be more readily viewed by reducing the “Change Selection Every” interval shown in [Figure 3-21](#) from 2 seconds to 1 second.

Rogue Access Points


Rogue access points are access points that are detected by the wireless LAN infrastructure and determined not to be members of the same RF group or WLAN system. In addition, any devices that are participating as members of ad-hoc networks are also detected as rogue access points (but with a rogue type of AD_HOC, unless location appliance rogue access point polling has been configured to exclude ad-hoc rogues).

Rogue access points are indicated on WCS location floor maps using an icon representing a skull-and-crossbones within a black circle, as shown in [Figure 3-22](#). They may be totally wireless, connected to the same wired infrastructure as the detecting WLAN, or connected to an entirely different wired infrastructure. To display rogue access points on the WCS location floor map, ensure that the **Rogue APs** checkbox option is enabled from the **Layers** dropdown selector at the top of the floor map display, and click **Load** in the left-hand column. To avoid excessive clutter, WCS will display the first 250 rogue access points on the floor map. To view the location of rogue access points beyond the first 250, rogue access point filtering must be used.

Figure 3-22 Rogue Access Point Location Map



It is possible to filter the location information displayed by the WCS based on the age of the information. In [Figure 3-22](#) WCS displays location appliance information that has aged up to 15 minutes. Alternatively, this value could be set to 2 or 5 minutes for more recent location information or ½, 1, 3, 6, 12, or 24 hours for older information.

By clicking on the blue chevron  that is displayed to the right of the **Rogue APs** checkbox, rogue access point filtering options can be specified and additional information can be displayed, such as:

- The total number of rogue access points detected on this floor.
- Small icons (shown above) or standard size icons can be selected. When using small icons, text is not displayed on the floor map for the rogue access point except when a mouse-over is performed. When using standard size icons, an on-screen tag displaying the MAC address of the rogue access point appears.
- Either all rogue access points can be displayed, or filtering can be performed to select which rogue access points to display on the floor map. This is based primarily on MAC address but can be augmented by filtering on the state of the rogue detection (Alert, Known, Acknowledged, Contained, Threat, or Known Contained) as well as whether or not the rogue access point was seen to be connected to the same wired network as the detecting wireless system. As mentioned previously, only up to 250 rogue access points will be shown at any one time on floor maps. If there are greater than 250 rogue access points detected, the total number found will be indicated in the left hand column status area during each communication cycle between WCS and the location appliance. It is recommended that filtering be used to reduce the total number of rogue access points selected for display if you receive this warning.

In software Release 4.1 of the Cisco UWN, WLAN controllers provide support for the maximum number of rogue access points shown in [Table 3-3](#).

Table 3-3 Maximum WLC Rogue Access Point Capacity

Controller Model	Rogue APs Supported
2006	125
2106	125
4402	625
4404	625
WiSM	1250
NM-WLC6	125
NME-WLC8/12	125
3750G	625

Complete information on any displayed rogue access point can be obtained simply by left-clicking the cursor on the circular skull-and-crossbones icon representing the desired rogue access point on the floor map. Doing this yields a screen containing detailed information as shown in [Figure 3-23](#). Note however, there is no RSSI information displayed for rogue access points when the location map is enlarged. Using the dropdown menu located in the upper right-hand corner, location history and playback information for the rogue access point in question can be accessed, similar in format and function to that described previously for WLAN clients and 802.11 active RFID tags.

Figure 3-23 Rogue Access Point Detailed Information

Logout | Refresh | Print View

Wireless Control System

Monitor ▾ Reports ▾ Configure ▾ Location ▾ Administration ▾ Help ▾

[Alarms](#) > Rogue - 00:1c:b0:eb:e2:30

General

Rogue MAC Address	00:1c:b0:eb:e2:30
Vendor	Unknown
Rogue Type	AP
On Network	No
Owner	
State	Alert
SSID	loc-wlc-04
Channel Number	6
Containment Level	Unassigned
Radio Type	b/g
Strongest AP RSSI	-77
No. of Rogue Clients	0
Created	Aug 15, 2007 3:23:13 PM
Modified	Aug 16, 2007 6:50:35 PM
Generated By	Controller
Severity	Minor
Previous Severity	Minor

Annotations

Annotations go here.

Message

Rogue AP '00:1c:b0:eb:e2:30' with SSID 'loc-wlc-04' and channel number '6' is detected by AP 'sjc14-32b-ap9' Radio type '802.11b' with RSSI '-77' and SNR '1'.

Help

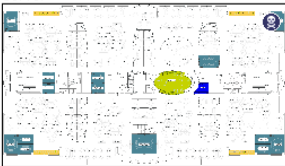
Rogue AP '00:1c:b0:eb:e2:30' with SSID 'loc-wlc-04' and channel number '6' is detected by AP 'sjc14-32b-ap9' Radio type '802.11b' with RSSI '-77' and SNR '1'.

Location Notifications

Absence	0
Containment	0
Distance	0
All	0

Location

Floor	Cisco S3 - Site 5>BLD 14>3rd floor
Last located at	Aug 16, 2007 6:59:12 PM
On Location Server	loc-2-2



[Enlarge](#)

[Rogue Clients](#)

[Event History](#)

-- Select a command --

GO

Set State to 'Unknown - Alert'

Set State to 'Known - Internal'

Set State to 'Acknowledged - External'

1 AP Containment

2 AP Containment

3 AP Containment

4 AP Containment

Location History

223319

Rogue Clients

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, as shown in Figure 3-24. To display rogue clients on the WCS location floor map, ensure that the **Rogue Clients** checkbox option is enabled from the **Layers** dropdown selector at the top of the floor map display, and click **Load** in the left-hand column. To avoid excessive clutter, WCS will display the first 250 rogue clients on the floor map. To view the location of rogue clients beyond the first 250, rogue client filtering must be used.

Figure 3-24 Rogue Client Location Map



It is possible to filter the location information displayed by WCS based on the age of the information. In Figure 3-24 WCS displays location appliance information that has aged up to 15 minutes. Alternatively this value could be set to 2 or 5 minutes for more recent location information or ½, 1, 3, 6, 12 or 24 hours for older information.

By clicking on the blue chevron ▶ that is displayed to the right of the **Rogue Clients** checkbox, rogue client filtering options can be specified and additional information can be displayed, such as:

- The total number of rogue clients detected on this floor.
- Small icons (shown above) or standard sized icons can be selected. When using small icons, no text is displayed on the floor map for the rogue client except when a mouse-over is performed. When using standard size icons, an on-screen tag displays the rogue client's MAC address.

- Either all rogue clients can be displayed or filtering can be performed to select which rogue clients to display on the floor map. Filtering can be based on the MAC address of rogue access point to which it is believed the rogue client is associated or it can be based on the state of the rogue client (alert, contained or threat). As mentioned previously, only up to 250 rogue clients will be shown at any one time on floor maps. If there are greater than 250 rogue clients detected, the total number found will be indicated in the left hand column status area during each communication cycle between WCS and the location appliance. It is recommended that filtering be used to reduce the total number of rogue clients selected for display if you receive this warning.

In software Release 4.1 of the Cisco UWN, WLAN controllers provide support for the maximum number of rogue clients shown in [Figure 3-24](#).

Table 3-4 Maximum WLC Rogue Client Capacity

Controller Model	Rogue Clients
2006	100
2106	100
4402	500
4404	500
WiSM	1000
NM-WLC6	100
NME-WLC8/12	100
3750G	500

Complete information on any displayed rogue client can be obtained simply by left-clicking the cursor on the rectangular black skull-and-crossbones icon representing the desired rogue client on the floor map. This yields the screen shown in [Figure 3-25](#). However, RSSI information is not displayed for rogue access points when the location map is enlarged.

Using the dropdown menu located in the upper right-hand corner, you can access location history and playback information for the rogue client that is similar in format and function to that described previously for WLAN clients, active RFID tags and rogue access points.

Figure 3-25 Rogue Client Detailed Information

Wireless Control System

Rogue Client "00:16:6f:1b:4b:dc"

Client MAC Address	00:16:6f:1b:4b:dc
Number of detecting APs	2
First Heard	Fri Aug 17 01:44:47 2007
Last Heard	Fri Aug 17 01:53:13 2007
Rogue AP MAC Address	00:14:1b:b6:ed:c0
Status	Alert

Location

Floor	Cisco SJ - Site 5>BLD 14>3rd floor
Last located at	Aug 16, 2007 6:59:07 PM
On Location Server	loc-2-2

Location Notifications

Absence	<u>0</u>
Containment	<u>0</u>
Distance	<u>0</u>
All	<u>0</u>

Map (High Resolution)

Location History

2233321

It is important to understand how localization of rogue access points and clients differs from that of WLAN clients and asset tags. Recall from prior discussion that WLAN clients transmit probe requests periodically across multiple channels. Because infrastructure access points are spending the vast majority of their time on their assigned channels, these probe requests tend to be detected quickly and relayed to the controllers to which the access points are registered. Asset tags do not transmit probe requests but rather multicast tag messages on the channels for which the tags have been configured. These multicasts are quickly detected by infrastructure access points operating on these channels in the vicinity of the asset tags.

Rogue devices may not be operating on the same channels to which your infrastructure access points have been assigned. Because of this, these rogue devices may be detected during periodic off-channel scans conducted by infrastructure access points. For an LWAPP access point operating in local mode, this off-channel scans typically occur for about 500 milliseconds out of every 180 seconds of operation (or about 50 milliseconds per non-primary channel per 180 second interval).

Workgroup Bridges

Some Cisco autonomous access points can connect to the Cisco UWN in a special mode of operation known as workgroup bridge (WGB) mode. Access points configured as workgroup bridges can provide wireless connectivity to the Cisco UWN for groups of wired clients, making the wired clients essentially appear as wireless clients to the UWN. Cisco AP1121, AP1130, AP1231, AP1240, and AP1310 access points containing Cisco IOS Release 12.4(3g)JA or greater (on 32-MB access points) or Cisco IOS Release 12.3(8)JEB or greater (on 16-MB access points) can be configured for workgroup bridge mode.

**Note**

For further information about the configuration of workgroup bridges and their role in the Cisco UWN, refer to *The Workgroup Bridge in a Lightweight Environment* located at the following URL: http://cisco.com/en/US/docs/wireless/access_point/12.4_3g_JA/configuration/guide/s43hot.html#wp1059452.

The default roaming behavior for workgroup bridges is to delay active scanning for potential access point roam candidates until the WGB has lost its association. While such behavior may be perfectly acceptable when workgroup bridges are used in stationary applications, it can cause concern in mobile WGB applications (such as a mobile cart-based array of Ethernet-only medical equipment) because of the following:

- Delaying the search for potential access point roam candidates until association is lost can introduce unnecessary application delays, which may negatively the performance of mobile timing-sensitive applications and cause application lockups or time-outs.
- Depending on the environment, mobile workgroup bridges may move about for considerable distances while associated to the same access point. In this case, the default WGB behavior will result in an absence of probe requests, causing the location appliance to rely on stale probe request RSSI information and potentially leading to poor WGB location fidelity until the WGB is faced with a roaming event.

Access points configured in WGB mode do not respond to broadcast Radio Measurement Requests that are sent as a result of the Cisco Compatible Extensions Location Measurement parameter being enabled. Therefore, Cisco Compatible Extensions Location Measurement cannot be used as a mechanism with which to trigger consistent periodic probing in work group bridges.

The Cisco IOS CLI **mobile station** command can be used on the workgroup bridge to provide a significant degree of improvement in workgroup bridge location fidelity. When you enable this setting in the workgroup bridge, it causes it to perform an active scan when it detects low access point RSSI, excessive radio interference, or a high percentage of frame loss. The workgroup bridge will use the information it learns from the active scan to determine whether any access points offering better service are available to it, and will roam to a new access point before it loses its current association.

The basic format of the command is:

```
mobile station period <seconds> threshold <|dBm|>
```

where the value for **period** denotes how often the workgroup bridge checks the RSSI of its currently associated access point, and the value for **threshold** specifies the absolute value of the minimum acceptable access point RSSI in dBm. The default values are 20 seconds and 70 dBm respectively.

**Note**

Complete details regarding the configuration of the Cisco IOS **mobile station** command can be found in *Cisco IOS 12.4(3g)JA for Access Points and Bridges*, “mobile station” command reference page located at the following URL:

http://cisco.com/en/US/docs/wireless/access_point/12.4_3g_JA/command/reference/cr43main.html#wp2593116

The values for period and threshold should be adjusted for your specific environment in order to balance the need for consistent and regular probe requests against the possibility of excessive roaming.

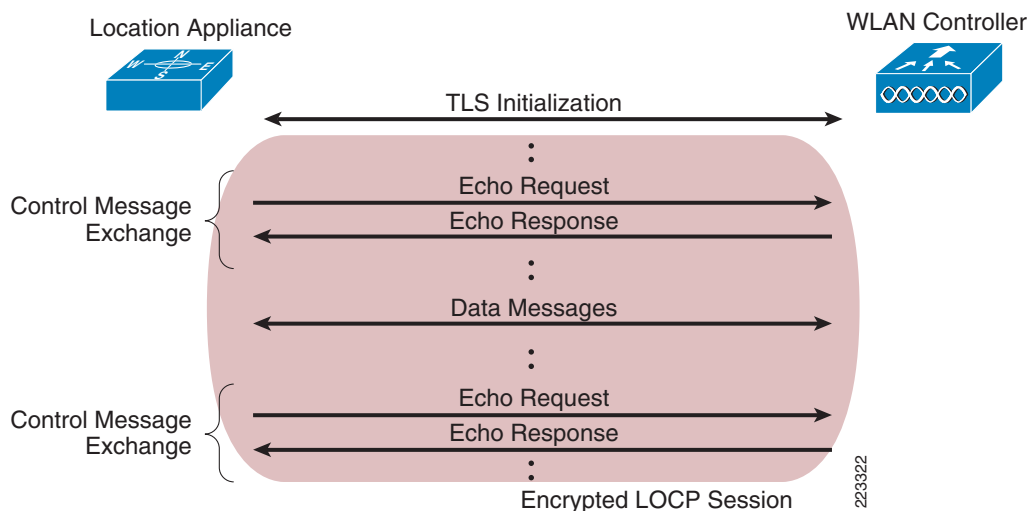
Decreasing the threshold value to a very low value, causing an active scan to always occur at each period interval, for example, will typically improve the location fidelity of work group bridges that seldom roam significantly. The trade-off with doing this however, is that such settings may also increase the frequency with which the workgroup bridge roams. However, this trade-off is generally viewed as equitable since in properly deployed environments with good coverage and access point placement, the increase in WGB roaming should be negligible whereas the improvement in mobile workgroup bridge location fidelity in cases where there is seldom roaming between access points can be very significant

Cisco Location Control Protocol (LOCP)

Cisco Unified Wireless Network (CUWN) software Release 4.1 introduces the Cisco Location Control Protocol (LOCP), an architectural enhancement that improves communication efficiency and supports new capabilities between the location appliances and one or more WLAN controllers. LOCP is a bi-directional protocol that can be run over a connection-oriented or connectionless transport. LOCP provides for an ongoing exchange of control messages that allows either endpoint to determine if its partner is still active.

Figure 3-26 illustrates the basic LOCP packet flow between the location appliance and each WLAN controller.

Figure 3-26 Location Appliance WLAN Controller LOCP Session




In Release 4.1 the location appliance is pre-configured by the user with regard to the IP addresses of the controllers it is to communicate with. Once the connection between the controller and the location appliance is initialized, an encrypted TLS session is established between the two endpoints over which

all further LOCP traffic will travel. Each endpoint periodically verifies that its partner is active and ready to accept requests by participating in Echo Request/Response control message exchange, as shown in Figure 3-26.

If a location appliance detects that a WLC is no longer responding, it will temporarily disable any other requests to that WLC until the WLC becomes active again. The connection between the location appliance and the WLC can be maintained for multiple exchanges (the typical case) or can be initiated and disconnected for each data request.

The basic parameters controlling the LOCP session between the location appliance and the WLAN controllers defined to it are specified using the Location Servers > LOCP Parameters panel, as shown in Figure 3-27.

Figure 3-27 Setting LOCP Session Parameters



The meaning of the LOCP session parameters shown in Figure 3-27 are as follows:

- *Echo Interval*—The minimum time interval, in seconds, between echo requests sent from the location appliance to the WLAN controller. Valid values are 1 to 120 seconds, with the default value being 15 seconds.
- *Neighbor Dead Interval*—The minimum time interval, in seconds, that the location appliance will wait before marking a WLAN controller not responding to its Echo Requests as “dead”. This value should not be less than twice the Echo Interval. Recommended values are 2 to 240 seconds, with the default value being 30 seconds.
- *Response Timeout*—The maximum time interval, in seconds, within which the WLAN controller must respond to requests sent by the location appliance. Valid values are 1 to 99,999 seconds, with the default value being 1 second.
- *Retransmit Interval*—The minimum time interval, in seconds, the location appliance waits before retransmitting a LOCP request when it does not get a response back from the WLAN controller. Valid values are 1 to 99,999 seconds, with the default value being 3 seconds.
- *Maximum Retransmits*—The maximum number of retransmissions that will be attempted by the location appliance when a response is not received for a LOCP Request. Valid values are 1 to 99,999 attempts, with the default value being 5 attempts.

Note that the first two parameters are applicable only to Echo Request and Echo Reply control messages while the remaining parameters pertain to all data messages (such as Information and Measurement Requests and Responses).

Cisco Unified Wireless Network software Release 4.1 introduces the initial phase of LOCP. In this release, LOCP is used to augment traditional SNMP polling by transporting Cisco Compatible Tag Extensions for Wi-Fi tags telemetry and notification traffic from WLAN controllers to the location appliance. This traffic includes:

- Cisco Compatible Extensions for Wi-Fi tag telemetry, such as:
 - Motion, temperature, pressure, humidity, distance, quantity and status.
 - Battery state and predicted remaining life.
- High priority Cisco Compatible Extensions tag traffic, such as:
 - Call button, tag detached and tamper alert events.
 - Entry into the range of a chokepoint trigger.
 - Vendor-specific tag information used by third party location clients.

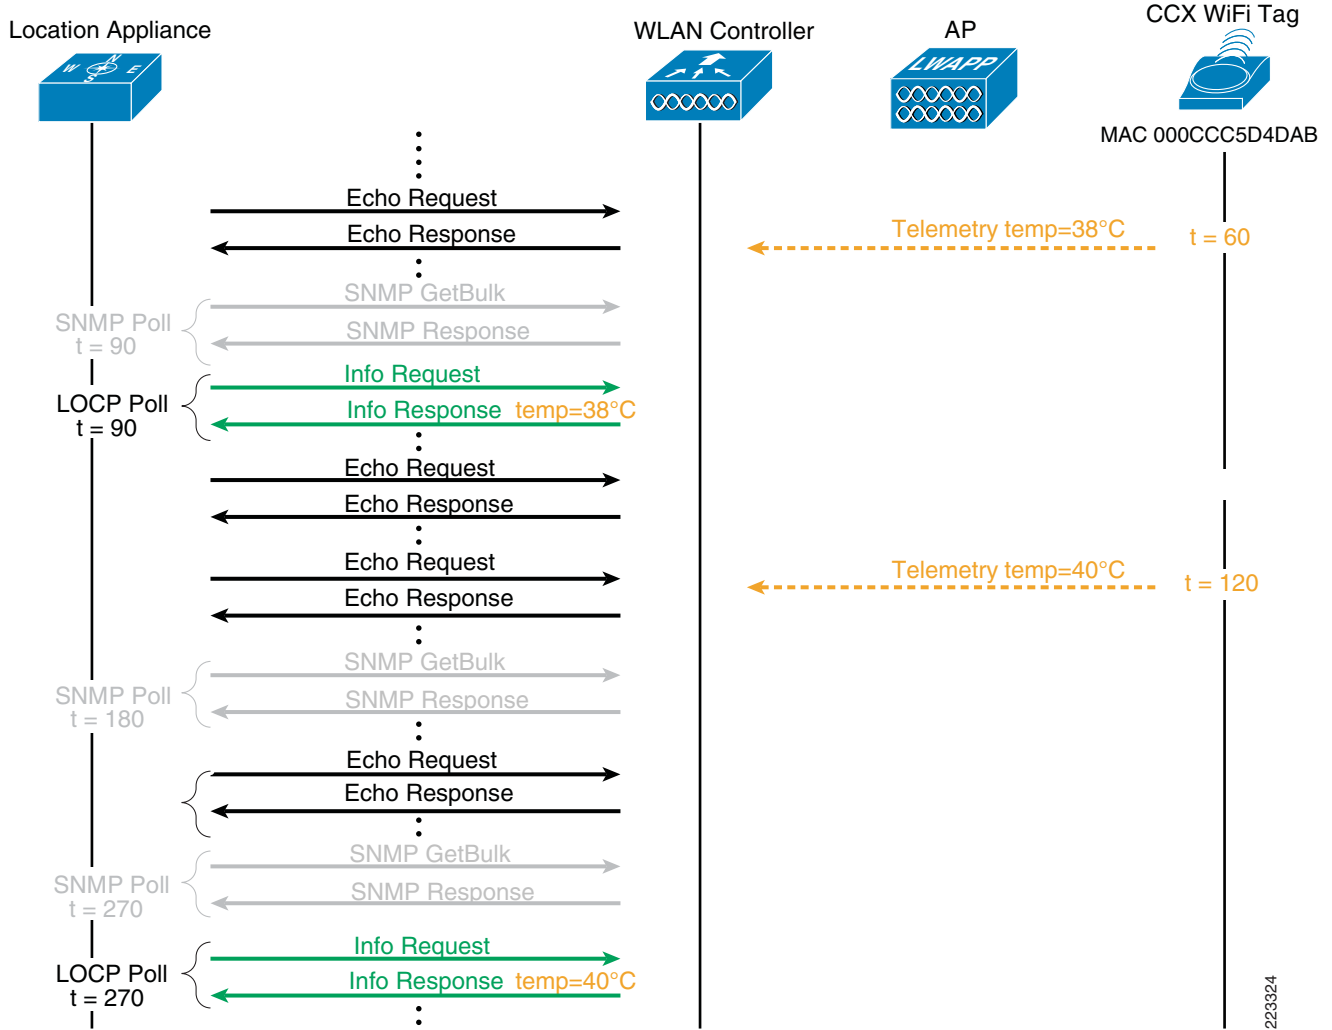
**Note**

Commencing with software Release 4.2, the Location Control Protocol (LOCP) receives additional enhancements and evolves into the Network Mobility Services Protocol (NMSP).

Asset Tag Telemetry Using LOCP

Beginning with Cisco UWN software Release 4.1, Wi-Fi RFID tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification may optionally pass tag telemetry information to the location-aware Cisco UWN as part of their tag message payload. This telemetry information is received by access points and collected by WLAN controllers. The location appliance periodically polls the WLAN controllers for tag telemetry using LOCP Information Requests. The controller will respond with the telemetry information it has received for each tag MAC address since the last LOCP polling cycle via a LOCP Information Response frame. These frames (as well as the polling exchange process) are illustrated in [Figure 3-28](#). Keep in mind that all frames shown between the location appliance and the WLAN controller travel are encrypted.

Figure 3-28 LOCP Information Request Polling



As you may have noticed in Figure 3-27, the LOCP polling interval is not directly configured in the location appliance configuration. Rather, it is derived from the asset tag SNMP polling interval occurring between the location appliance and WLAN controllers. In order to determine when a LOCP poll should be sent, the location appliance evaluates the following conditions during each controller polling cycle:

- Whether the controller software release supports LOCP. LOCP polls will not be sent to controllers that are not LOCP-capable.
- If the time interval since the last LOCP poll is 180 seconds, then LOCP polling will be performed during this asset tag SNMP polling cycle. LOCP Information Requests will be sent to all LOCP-capable controllers currently defined to the location appliance.
- If the interval since the last LOCP poll < 180 seconds, then LOCP polling will not be performed during this asset tag SNMP polling cycle.

The LOCP polling interval used by a location appliance to collect asset tag telemetry information in software Release 4.1 can be calculated from the asset tag SNMP polling interval using the following formula¹:

$$Poll_{LOCP} = \left\lceil \frac{180}{Poll_{TAG}} \right\rceil * Poll_{TAG}$$

$Poll_{LOCP}$ represents the LOCP polling interval and $Poll_{TAG}$ specifies the poll interval at which the location appliance polls the controller for asset tag location information via SNMP. Both of these values are specified in seconds. The value for $Poll_{TAG}$ is configured in WCS using Location Servers > Polling Parameters. For example, using an asset tag SNMP polling interval ($Poll_{TAG}$) of 120 seconds, the LOCP polling interval ($Poll_{LOCP}$) used by the location appliance is calculated to be 240 seconds.

Figure 3-28 helps to provide clarity to understanding the relationship between LOCP and SNMP polls and their impact on the receipt of tag telemetry. We see that temperature telemetry is transmitted from a Cisco Compatible Extensions for Wi-Fi Tags compatible tag with MAC address 00:0C:CC:5D:4D:AB at time $t=60$ seconds. This transmission is received by one or more access points. These access points pass the telemetry information (temperature of 38°C in our example) to their respective registered WLAN controllers. Since Figure 3-28 shows that the SNMP poll occurring at time $t=90$ seconds is the very first poll (which also implies that no previous LOCP polls have occurred), the controller will receive a LOCP poll at time $t=90$ seconds as well. This is indicated in the figure by the receipt of the Information Request frame. The controller responds to the poll by passing any accumulated tag telemetry information to the location appliance in a LOCP Information Response frame.

If tags are configured to send multiple frame copies (or bursts) per channel, the controller eliminates any duplicate tag telemetry and passes the distilled telemetry values to the location appliance. The location appliance then updates its databases with this telemetry information and makes it available to location clients via the SOAP/XML API.

Subsequent inbound telemetry is handled in an analogous fashion. For example, in Figure 3-28 at time $t=120$ seconds we see an inbound temperature telemetry update indicating that the temperature has increased to 40 °C. The aforementioned cycle of events would reoccur, culminating in the telemetry update being transmitted to the location appliance at time $t=270$ seconds. It is important to understand why the tag telemetry is passed to the location appliance at time $t=270$ seconds and not at time $t=180$ seconds, which is where we observe an SNMP poll occurring. As explained earlier, LOCP polling only occurs if the time delta since the last LOCP poll is 180 seconds or more. At time $t=180$ seconds, the time delta since the previous LOCP poll is only 90 seconds, thus no LOCP poll occurs at that time.

While Figure 3-28 illustrates the simple case of a single tag passing only a single telemetry value, it should be noted that LOCP is designed to efficiently transport telemetry values from multiple tags just as easily. Each Information Response frame allows multiple tag MAC addresses to be specified by the controller, with each MAC address being associated with one or more telemetry values. For example, instead of passing only temperature telemetry, the tag shown in Figure 3-28 could include temperature, pressure, humidity and so on. All of this information would be included in the Information Response transmitted at the next LOCP poll. Inbound telemetry traffic from multiple tags would be aggregated by the controller in a similar fashion, with each LOCP endpoint capable of performing LOCP frame fragmentation and reassembly if necessary.

With the exception of battery state information which can be “pushed” via asynchronous northbound notifications from the location appliance, tag telemetry is made available to location clients only via the SOAP/XML API.


1. The ceiling function “ $\lceil x \rceil$,” represents the application of the ceiling function to the positive integer x . This rounds up x upwards, returning the smallest integer that is greater than or equal to x .

For information regarding tag telemetry deployment considerations, refer to [Tag Telemetry and Notification Considerations](#), page 6-27.

Asset Tag Notifications Using LOCP

Beginning with Cisco UWN software Release 4.1, Wi-Fi active RFID tags compliant with the Cisco Compatible Extensions Wi-Fi tag specification can pass optional high priority, chokepoint and vendor-specific notification events to the location-aware Cisco UWN. Indication of high-priority tag events are received by one or more access points via tag multicast messages frames that contain additional payload information indicating the nature of the event. This information is typically dispatched by asset tags at the time the tag detects that the event has occurred. Once received by the WLAN controller, these time-critical events are handled outside the polled LOCP polling process and passed immediately to the location appliance using a LOCP Notification frame, as shown in [Figure 3-29](#).

Figure 3-29 LOCP Notifications



[Figure 3-29](#) shows the two basic types of LOCP notifications supported in Cisco UWN software Release 4.1, the *Measurement Notification* and the *Information Notification*:

- Measurement Notifications**—In Release 4.1, measurement notifications are used to convey information regarding the identity of any chokepoint proximity devices into whose range a tag may have entered. Tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification include chokepoint identification in the tag content field (such as the chokepoint MAC address). This information is passed to the location appliance along with the tag MAC address in real time, and can be used by location clients (such as WCS) to indicate that a tag is now within range (or out of range) of a particular chokepoint.

- *Information Notifications*—In Release 4.1, information notifications are used to convey vendor-specific data and tag high-priority events, such as:
 - When a tag user depresses a tag call button
 - When a tag detects that it has been removed from its carrier or attached asset
 - When a tag detects tampering.
 - When any other high-priority tag events occur.

While each tag vendor is responsible for determining the precise set of capabilities they choose to include in their product offering, the Cisco Compatible Extensions for Wi-Fi Tags specification provides for high-priority information to be uniformly included in the tag content field. This information is passed to the location appliance along with the tag MAC address, and can be used by location clients (such as WCS) to indicate that a high-priority tag event has taken place.

The Cisco Compatible Extensions for Wi-Fi Tags specification allows each tag vendor to pass vendor-specific information (such as proprietary tag messages or additional vendor-specific chokepoint information) from their tags into the Cisco UWN in real time. Vendor-specific information will be made available unaltered to location clients via the SOAP/XML API interface of the Cisco location appliance. Vendor-specific information that is sent within a high-priority event can also be “pushed” to location clients from the location appliance via an asynchronous northbound notification from the location appliance to location clients using SMTP, UDP-Syslog, SNMP traps or SOAP transports. Location clients must be capable of receiving and processing these notifications on the aforementioned ports in order to provide real-time notification of such events to end users.

[Figure 3-29](#) clearly indicates that unlike tag telemetry, there is no dependency on any polling mechanism between the location appliance and the WLAN controller. LOCP notifications will be generated from the WLAN controller to the location appliance as tags communicate those events. High-priority tag events, vendor specific data and chokepoint in-range information are not aggregated by WLAN controllers in Release 4.1 of the Cisco UWN. Each incoming tag multicast message bearing such information, received by a WLAN controller from each registered access point, results in the generation of an information notification or a measurement notification frame from the WLAN controller to the location appliance.

No dependency exists between the transmission of tag telemetry and the transmission of chokepoint, high-priority or vendor specific information. For example, a tag may be relaying telemetry information about an asset as it traverses through the range of chokepoints in an environment. As seen in [Figure 3-29](#), the tag telemetry is retained by the WLAN controller until the next LOCP polling interval, whereas indication of chokepoints that are in-range is sent immediately by the WLAN controller to the location appliance in the form of a measurement notification.

In addition to providing updated information to clients via the SOAP/XMP API (for example, when a location client issues a XML GetTagInfo or GetTagLocation request), the location appliance can also dispatch external asynchronous northbound notifications upon receipt of LOCP measurement or information notifications (refer once again to [Figure 3-29](#)). These external northbound notifications are sent for the following conditions using SNMP, SMTP, SOAP or UDP-Syslog transports:

- Call Button, Tag Tampering, or Tag Detached
- Chokepoint in-range
- Other Priority Events (user defined)

For information regarding deployment considerations surrounding tag notifications, refer to [Tag Telemetry and Notification Considerations, page 6-27](#).



CHAPTER 4

Installation and Configuration

Installing and Configuring the Location Appliance

Detailed procedures for installing and configuring the Cisco Location Appliance can be found in the following documents:

- *Release Notes for Cisco Wireless Location Appliance Release 3.0*—
http://www.cisco.com/en/US/products/ps6386/prod_release_notes_list.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_installation_and_configuration_guides_list.html

Configuring Cisco WCS for Location Tracking

It is assumed that the reader has installed either a Windows or Linux-based version of WCS that is appropriately licensed for location use with the Cisco Wireless Location Appliance. Detailed procedures for configuring the Wireless Control System for location use with the Cisco Wireless Location Appliance can be found in the following documents:

- *Cisco Wireless Control System Release Notes, Release 4.1*—
http://www.cisco.com/en/US/products/ps6305/prod_release_notes_list.html
- *Cisco Wireless Control System Configuration Guide, Release 4.1*—
http://www.cisco.com/en/US/products/ps6305/products_installation_and_configuration_guides_list.html
- *Cisco Wireless Location Appliance: Deployment Guide*—
http://www.cisco.com/en/US/products/ps6386/prod_technical_reference09186a008059ce31.html

Configuring Location Appliance History Parameters

The configuration of Location Server > Administration > History Parameters is discussed in the document entitled *Cisco Wireless Location Appliance Configuration Guide: Editing History Parameters* at the following URL:

http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a008082d72f.html#wp1046373.

Further clarification regarding some of these parameters is provided in the subsections that follow.

A common misconception about the history capabilities of the location appliance is that it somehow stores a historical record of all locations the client has ever encountered. As is discussed in the following two sections, the location application stores history information based on the values of the archive period and archive interval parameters. If a history record for a device is recorded at time T_0 and the archive period is 30, the next history record for that device is written at T_{0+30} . The device may have undergone several changes in location between T_0 and T_{0+30} ; however, only the location states at time T_0 and T_{0+30} are recorded in the history database.

History Archive Period

The history archive period (shown as “Archive For”) specifies the number of days that the location appliance retains location history records for each enabled history collection category. The default archive period is 30 days. Changes to the default history archive period should be done with careful consideration after consultation with your Cisco field technical representative or the Cisco Technical Assistance Center, because longer history periods typically increase the amount of space consumed by the location history database. Because newer history data within the archive period does not overwrite older data, the combination of a large number of devices, an injudicious selection of history categories, and an excessive history archive period can increase the risk of exhausting available free space.

To illustrate this point, we can compare the amount of disk storage that is consumed when selecting one combination of history category, archival period, and archival interval versus another. To do this, let us assume an environment consisting of 1100 WLAN clients, 300 asset tags, 20 rogue access points, and 30 rogue clients. [Figure 4-1](#) illustrates the effect on consumed disk storage of the following:

- Increasing the default archive period to 365 days for all device categories
- Reducing the default history archive interval for clients (“mobile devices”) and asset tags to 60 minutes.

Figure 4-1 Impact of History Interval and Archive Period on Database Size

		<u>Location History</u> <u>(bytes)</u>			<u>Location History</u> <u>bytes</u>
Number of Mobile Devices =	1100	28,248,000	Number of Mobile Devices =	1100	2,062,104,000
History Interval =	360 mins		History Interval =	60 mins	
Archive Period =	30 days		Archive Period =	365 days	
Number of Tags =	300	1,296,000	Number of Tags =	300	189,216,000
History Interval =	720 mins		History Interval =	60 mins	
Archive Period =	30 days		Archive Period =	365 days	
Number of Rogue APs =	20	117,600	Number of Rogue APs =	20	1,430,800
History Interval =	720 mins		History Interval =	720 mins	
Archive Period =	30 days		Archive Period =	365 days	
Number of Rogue Clients =	30	118,800	Number of Rogue Clients =	30	1,445,400
History Interval =	720 mins		History Interval =	720 mins	
Archive Period =	30 days		Archive Period =	365 days	
		29,780,400 bytes			2,254,196,200 bytes

190558

Although the estimates shown in Figure 4-1 are only an approximation (they do not account for per record display string sizes and database overhead, for example), you can see that database size increases from about 30 MB to over 2.25 GB because of these changes in location history alone. The database backup mechanism on the location appliance requires that there be at least as much free space available as is used in order to support reliable extraction and compression, thereby bringing the total estimated space requirement to over 5 GB.

History Database Pruning

Database pruning is especially important in situations when there is a high risk of a situation occurring where available hard disk space becomes critically low. If low available disk space situations re-occur, more aggressive data pruning intervals may be warranted such that pruning occurs more frequently and well in advance of a low disk space situation. These aggressive data pruning intervals may need to be combined with a shorter history archive interval if the adjusted pruning intervals alone are not sufficient in addressing the low free disk space situation.

Configuring Location Appliance Advanced Parameters

The configuration of Location Server > Administration > Advanced Parameters is discussed in the document entitled *Cisco Wireless Location Appliance Configuration Guide: Editing Advanced Parameters* at the following URL:

http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a008082d72f.html#wp1050981.

Further clarification regarding a subset of the Advanced Parameters is provided in the following subsections.

Absent Data Cleanup Interval

The “Absent Data Cleanup Interval” or ADCI (range 1 to 99,999 minutes) specifies the amount of time that an entry is kept for a tracked entity (WLAN client, tag, rogue access point, or rogue client) in the active location database. The ADCI specifies the amount of time that must expire before the tracked device entry is removed from the active location database if no recent updates have been received for that device.

For example, if the RSSI information for an asset tag was last recorded by the location appliance two days ago and the cleanup interval is set to the default value of 1440 minutes (24 hours or 1 day), the station will be removed from the active location database after the expiration of the 24 hour absent data cleanup interval. Note that once the device is removed from the active location database, it will not be possible to “scroll back” and review the last known location of the device using the “load location server data as old as” dropdown menu control.

The limit of 2500 total tracked devices in the location appliance applies strictly to those devices that are in the active location database. Once the total number of devices (clients, tags and rogues) in the active database reaches 2500, additional devices cannot be tracked by this location appliance until some of the currently tracked devices contained in the active location database expire and are pruned from the database.

In some cases, the default value for the Absent Data Cleanup Interval may be found to excessively delay the clean-up of devices that have been recently removed from the tracked environment. A good case in point might be a location appliance with 1500 tracked asset tags and client stations, where an operator has enabled rogue location tracking in an environment with a high concentration of rogue devices. If the system were to discover 1000 rogue devices, for example, these would be added to the active location database and would bring the total number of tracked devices for this location appliance to its maximum capacity of 2500. If location tracking of tagged assets and WLAN clients are considered to be a higher priority than the tracking of rogue devices, a potential problem could exist. If new tags or clients are added to the environment, there may not be any available capacity to track them until some of the currently tracked devices (existing WLAN clients, asset tags or rogue devices) expire and are pruned from the active location database.

**Note**

Version 4.2 of the location-aware Cisco UWN introduces an enhancement that allows for individual limits to be placed on what portion of the location appliance’s aggregate tracked device capacity is allocated to each tracked device category (i.e. WLAN clients, asset tags, or rogue devices).

It is important to note that in this situation, disabling the tracking of rogue devices in the location appliance entirely will not immediately remove the 1,000 tracked but unwanted rogue devices from the active location database. Rather, the Absent Data Cleanup Interval will by default maintain each currently tracked rogue device in the active location database for a period of 1440 minutes past the time

of its last RSSI update. Plainly put, in the case of our example using the default value for the ADCI, disabling rogue location tracking today will not prune those tracked device entries from the active location database until approximately the same time tomorrow.

To work around this and remove the unwanted devices from the active location database more expeditiously, we can temporarily set the Absent Data Cleanup Interval to a much lower value for a brief duration in order to accelerate the pruning of any unwanted tracked devices from the active location database. For example, our hypothetical operator might choose to temporarily set the Absent Data Cleanup Interval to sixty minutes after disabling location appliance polling for rogue devices. Sixty minutes after this setting has been applied, the location appliance will remove all devices from the active location database for which updated information has not been received from WLAN controllers within the last hour, including the undesired rogue devices. Once this has occurred, the “Number of Tracked Elements” field shown on the Location Servers > Advanced Parameters menu page should decrease, reflecting the number of devices removed from the active location database.

The Absent Data Cleanup Interval is a single parameter that applies to all device categories. Thus, a potential drawback of temporarily lowering the ADCI in this way is that the removal of tracked devices from the active location database occurs in a non-selective fashion. That is to say, all devices for which information updates have not been received by the location appliance meeting the ADCI time criteria will be removed. In our example, this means that not only would our unwanted rogues be removed, but so would any clients or asset tags for whom we have not received any updates in the last sixty minutes as well. Such behavior could prove surprising to location client users that have come to depend on a 24 hour window of prior location information in order to locate “lost” assets for which current location information is not available.

In lab testing, it was found that the use of the location history database can partially mitigate this in cases where tracked devices have been removed from the active location database but are later re-detected by the UWN. In such cases, any prior collected location history records will once again be available, unless their history archive period has expired and the history records themselves have been pruned. History records for devices that have been deleted from the active location database, but have not been re-detected by access points, will not be accessible via location clients.

In some cases, it may be desirable to alter the default value for Absent Data Cleanup Interval on a more permanent basis. One example of such a case might be an facility that employs location tracking but has a large number of transient Wi-Fi devices, such as client laptops, PDAs and so on, residing onsite only a few hours before moving on, and then not returning for several days. Another example might be a logistics cross-docking facility that may only contain 2000 tagged asset containers during any four hour period, but through whose doors a volume of 10,000 or more tagged asset containers may pass within a 24 hour period.

In either of these cases, the quantity of track-able Wi-Fi client devices or asset tags actually on site at any one time may be significantly less than the maximum tracked device capacity of the location appliance. However, the number of transient devices that may pass through the facility over a 24 hour period will easily exceed 2500. Should this occur while using the default value for Absent Data Cleanup Interval, there may be a risk of the location appliance's tracking capacity becoming exhausted, as devices that may have left the facility several hours ago will not be removed from the active location database until the 24 hour ADCI has expired. Setting the value for Absent Data Cleanup Interval to a lower value (say, for example, four hours or 240 minutes) would expedite the cleanup of these migrated devices and release tracked device capacity on the location appliance for use by recent device arrivals.

Reducing the value of the Absent Data Cleanup Interval is not without its tradeoffs, however. For further discussion of the Absent Data Cleanup Interval, the potential tradeoffs involved in changing it and how this may factor into your overall design approach, it is recommended that the reader consult the examples given in [Multiple Location Appliance Designs, page 5-32](#).

Memory Information

- **DB Disk Memory**—A misnomer, this parameter does not refer to “memory” on the location appliance. Rather, it displays the amount of disk space that has been consumed by the location appliance database. This information is useful when determining whether a database de-fragmentation should be performed (see [Advanced Commands, page 4-6](#)).
- **Run Java GC**—This command runs a general memory clean-up immediately. Normally, memory cleanup is initiated by the system automatically and thus does not require manual initiation. Therefore, Java General Cleanup need only be run when directed by the Cisco Technical Assistance Center (TAC) or Cisco Engineering.

Advanced Commands

The Defragment Database advanced command defragments the location database and reclaims allocated but unused disk space. A database defragmentation can be beneficial if free disk space on the location appliance is running low because of large database size, or if the response time of the location appliance appears noticeably slower when data is requested by location clients.

To determine how much free space is currently available on the location appliance, it is necessary to log into the location appliance via either the CLI serial console or an SSH session. When logged in, use the Linux command **df -H** to display disk free space, as follows:

```
[root@AeS_Loc root]# df -H

Filesystem      Size  Used Avail Use% Mounted on
/dev/sda2       77G   3.2G   70G   5% /
/dev/sda1       104M   16M   83M   16% /boot
none            526M    0   526M   0% /dev/shm

[root@AeS_Loc root]#
```



Note

The **df -H** command is used above because it is a commonplace practice for most computer disk manufacturers to assume 1 GB = 1,000,000,000 bytes. The **-H** option displays output as powers of 1000 rather than 1024. Use **df -h** if your preference is for the contrary.

The **df** display output shown here is for a location appliance containing a hard disk drive with an unformatted capacity of 80 GB. Notice that there are two main file systems defined: `/dev/sda1`, which is the Linux boot file system; and `/dev/sda2`, which contains the root directory as well as the location application and all databases. You can clearly see from the display above that only 5 percent of all available space on `/dev/sda2` is currently being used. That being the case, there is an abundance of free space available and defragmentation is unlikely to be required at this time.

You can use the information in the **df** output along with the knowledge of the size of the location database (from DB Disk Memory described in [Memory Information, page 4-6](#)) to approximate the maximum recommended size to which the location appliance database should be allowed to grow. At first glance, this may appear intuitive; that is, max recommended database size = total available disk space – (OS size + location application size). However, you should also account for the creation of a flat file that is used during the database backup process. Using the formula below, you can calculate the maximum recommended size of the location database including this additional free space plus a small additional amount to account for system overhead (such as the downloading of an location appliance upgrade image):

$$MaxDatabaseSize = \frac{TotalSpace - OSAppSpace}{2.3}$$

Where:

- *MaxDatabaseSize* is the maximum recommended size of the database in bytes



Note *MaxDatabaseSize* assumes the user has performed a cleanup of any residual location appliance upgrade images. Multiple residual upgrade images may consume additional free space exceeding these allotments.

- *TotalSpace* is the total amount of available space on /dev/sda2 in GB.
- *OSAppSpace* is the amount of space occupied by the Linux OS and the location appliance application on /dev/sda2. This can be calculated for the example shown above as:
(the amount of used disk space in Gigabytes) – (the current size of the location appliance database in Gigabytes).

The current size of the location appliance database can be found at WCS > Location Server > Advanced Parameters > DB Disk Memory. In the case of the system used for this example, DB Disk Memory = 24,608,768 bytes or .0246 GB. Thus, OSAppSpace = (3.2 - .0246 GB) or 3.175 GB.

Substituting the values for TotalSpace and OSAppSpace into the equation, you can calculate the maximum recommended size to which the location appliance database should be allowed to grow as (77 GB - 3.175 GB) / 2.3 = 73.825 / 2.3 = 32.0 GB. Therefore, to ensure proper operation of the database backup mechanism in a location appliance with an 80 GB unformatted capacity hard disk drive, the maximum recommended size of the location database (as indicated by DB Disk Memory) should not be allowed to exceed 32 GB.

Configuring Location Appliance Location Parameters

The configuration of Location Server > Administration > Location Parameters is discussed in *Cisco Wireless Location Appliance Configuration Guide: Editing Location Parameters* at the following URL: http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a008082d72f.html#wp1050973.

Further clarification regarding select parameters is provided in subsequent sections.

Enable Calculation Time

The *enable calculation time* location parameter refers to an advanced debugging option that enables logging of the amount of time that internal localization calculations consume. It is disabled by default and should be enabled *only* on the recommendation of the Cisco Technical Assistance Center (TAC) or Cisco Engineering, because it adds overhead to the location calculations.

Enable OW (Outer Wall) Location

Although the WCS Map Editor allows interior walls to be placed within floor maps, the location appliance only takes into consideration up to 50 “heavy” walls when evaluating path loss models and conducting positioning calculations. Heavy walls are those defined in the Map Editor with attenuation values of 13 dB. When the “Wall Usage Calibration” parameter in WCS (Monitor > Maps > Properties > Wall Usage Calibration) is set to “Auto”, the location appliance will dynamically determine whether to use the attenuation introduced by heavy walls during the calculations performed as part of the calibration process. The system administrator can, however, opt to include heavy wall attenuation in all cases by setting this parameter to “Use Walls”, or disable the use of heavy walls entirely by setting it to “Do Not Use Walls”.

Enable OW Location is a parameter that was used with software releases prior to release 4.0 of the Cisco UWN (i.e., prior to Release 2.1 of the Cisco Wireless Location Appliance). *Enable OW Location* is still displayed on the Location Server > Administration > Location Parameters menu in Release 4.1 of WCS for backward compatibility with these earlier releases. However, with software Release 4.1, there is no benefit to be gained by changing this parameter from its default setting.

RSSI Discard Times

- **Relative RSSI Discard Time**—This parameter denotes the relative boundary of RSSI sample times used in location calculations. It specifies the time between the most recent RSSI sample and the oldest usable RSSI sample. The default relative RSSI discard time is 3 minutes. During normal operation of the location appliance, this parameter should be left at the default value and should *not* be changed except on the advice and recommendation of the Cisco Technical Assistance Center (TAC) or Cisco Engineering.
- **Absolute RSSI Discard Time**—This parameter denotes the absolute boundary of RSSI sample times used in location calculations. The default is 60 minutes, which means that RSSI samples older than 60 minutes are not used in location calculations, regardless of relative RSSI discard time. During normal operation of the location appliance, this parameter should be left at the default value and should *not* be changed except on advice of the Cisco Technical Assistance Center (TAC) or Cisco Engineering.

RSSI Cutoff

In addition to enforcing the aforementioned relative and absolute time constraints against received RSSI reports, the location appliance also applies a parameter known as the *RSSI cutoff*. Subject to the time constraints described in *RSSI Discard Times*, the location appliance retains the four highest signal strength reports plus any signal strength reports that meet or exceed the value specified for RSSI cutoff. The default value for RSSI cutoff is -75 dBm.

The application of the RSSI cutoff threshold is illustrated in the following examples:

- Four RSSI reports of -68dBm, -70dBm, -72dBm, and -80dBm—All four reports are retained because they are the four highest reports.
- Five RSSI reports of -66dBm, -68dBm, -70dBm, -72dBm, and -74dBm—All five reports are retained because they all meet or exceed the default RSSI cutoff threshold.
- Five RSSI reports of -66dBm, -68dBm, -70dBm, -72dBm, and -80dBm—The first four reports are retained, the fifth report of -80dBm is discarded because it does not meet the default RSSI cutoff threshold of -75 dBm and there already exists four other signal reports that meet or exceed the threshold.

Configuring Location Appliance Notification Parameters

The configuration of Location Server > Administration > Notification Parameters is discussed in *Cisco Wireless Location Appliance Configuration Guide: Configuring Notification Parameters* at the following URL:

http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a008082d744.html#wp1053921.

Further clarification regarding select parameters is provided in the following sections.

Queue Limit

The Queue Limit parameter specifies the size of the output notification queue of the location appliance. This value normally defaults to 500. The location appliance drops any outbound notifications above this limit if the output notification queue size is exceeded. Therefore, if you notice that some outbound notifications are being dropped (via the Notifications Dropped field), you may want to increase the queue limit size.

Retry Count

For each matching condition, the Retry Count specifies the number of northbound notification “firings” that will be allowed for the same device (over and above the initial firing) before the wait period specified by the Refresh Time parameter begins. Thus, the total number of “firings” of northbound notifications allowed between Refresh Time periods will be equal to one plus the value specified for Retry Count. The default value for Retry Count is one.

Keep in mind that:

- More than one physical northbound notification message can be sent per “firing” (for example SMTP, Syslog, SNMP, or SOAP).
- Retry Count and Refresh Time apply independently to each matching device MAC address.
- Retry Count and Refresh Time apply independently to event definitions. However, event definitions that apply the same trigger conditions to the same device MAC addresses will share a Retry Count/Refresh Time parameter set.

As an example, assume that the location appliance has been configured to transmit SNMP, email and syslog northbound notifications when a high-priority condition arises for a specific asset tag. For the purpose of this example, let us assume the high priority event is the depression of a call button on an asset tag, such as the AeroScout T2 or T3. Also assume that the notification Refresh Time is set to sixty minutes and the notification Retry Count is set to one (the default values). Under these conditions, the location appliance will generate SNMP, email and syslog northbound notifications for each high-priority event occurring for this tag, up to a maximum of two northbound notifications. After the value of one plus the Retry Count has been reached, the location appliance will skip firing any further northbound notifications for this condition and device for the time period specified by the Refresh Time. Once the Refresh Time has expired, this cycle will repeat unless the event has been cleared.

Refresh Time

Refresh Time specifies the length of the wait period between transmission of northbound notification sets for a specific event condition and device, as described above in [Retry Count, page 4-9](#). After the expiration of the Refresh Time, the event condition is eligible for re-evaluation and, if still present, may once again result in the generation of northbound notifications.

Refresh Time and Retry Count are used cooperatively to help limit the number of northbound notifications that are repeatedly generated for uncleared events. Retry Count limits the number of northbound notifications that are sent by the location appliance, while Refresh Time imposes a “waiting period” during which time no further northbound notifications will be sent for this event condition and device.

Refresh Time is specified in minutes, with the default being 60 minutes.

Notifications Dropped

This is a read-only counter field indicating the total number of notifications that have been dropped from the notification queue since the location appliance was started. Note that stopping and restarting the location appliance software application (locserverd) will reset this counter. The Notifications Dropped counter should be used in conjunction with the Queue Limit parameter to reduce the number of total dropped notifications.

Configuring Location Appliance LOCP Parameters

The configuration of Location Server > Administration > LOCP Parameters is discussed in *Cisco Wireless Location Appliance Configuration Guide: Configuring LOCP Parameters* at the following URL:

http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a008082d72f.html#wp1050918.

Location Appliance Dual Ethernet Operation

The Cisco Wireless Location Appliance is equipped with two 10/100/1000BASE-T Gigabit Ethernet ports that can be used to “dual-home” the location appliance to two different IP networks. This makes it a simple affair, for example, to configure a location appliance for service on network “A” while affording it the capability to be managed out-of-band on network “B” if the need arises. Complete step-by-step guidelines to accomplish this are available in *Cisco Wireless Location Appliance Installation Guide: Configuring the Location Appliance* at the following URL:

http://www.cisco.com/en/US/products/ps6386/products_installation_and_configuration_guide_chapter09186a00804fab8e.html#wp1040488.

Particular attention should be paid to the fact that the dual onboard Ethernet controllers on the location appliance are *not* intended for redundant or simultaneous connection to the same IP network. Configurations aimed at establishing parallel, load balancing, or redundant Ethernet connections to the same IP network are not recommended at this time.

Changing Location Appliance Default Passwords

Changing the “root” User Linux System Password

The location appliance ships with a default root userid and password. It is recommended that the password for the root userid be changed during initial configuration of the location appliance to ensure optimum network security. This can be done during the execution of the initial setup script as described in “Installation and Configuration” section of the *Cisco 2700 Series Location Appliance Installation and Configuration Guide*, located at

<http://www.cisco.com/en/US/docs/wireless/location/2700/quick/guide/li31main.html#wp1049597>.

When logged in, the Linux command **passwd** can be used to change the root system password as follows:

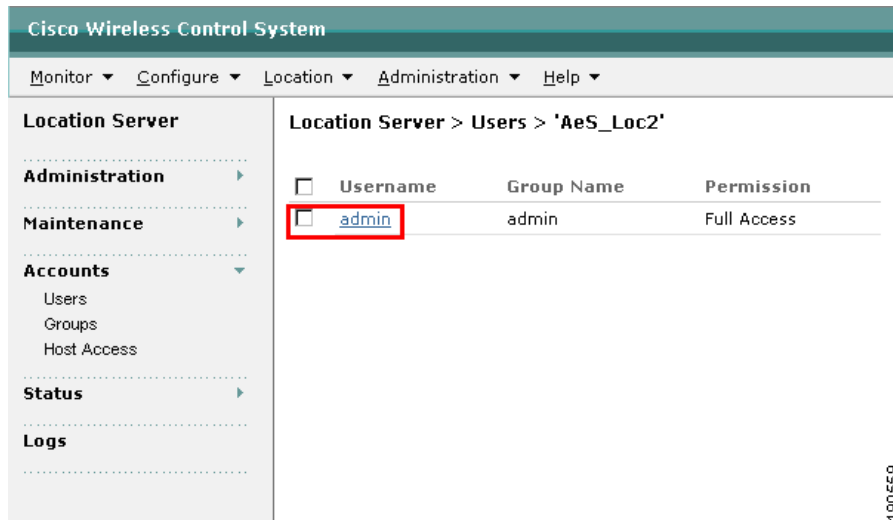
```
AeS_Loc login: root
Password:
Last login: Thu Oct 22 09:53:21 on ttyS0
[root@AeS_Loc root]# passwd
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
[root@AeS_Loc root]#
```

Changing the “admin” Location Server Application Password

The location server application on the location appliance ships with an administrator user account and group predefined. The userid is *admin* and the password is *admin*. After WCS has successfully contacted the location server application using the factory default administrator credentials, the default password on the admin account can be changed to a less well-known value via the WCS menu Location > Accounts > Users menu.

Step 1 Begin by clicking on the **admin** userid, as shown in [Figure 4-2](#).

Figure 4-2 Default Location Appliance User ID



Step 2 Clicking on the Admin box brings up the menu shown in [Figure 4-3](#), which allows the password to be changed for the admin userid.

Figure 4-3 Modifying the Admin Password

- Step 3** Finally, change the value for the password used by WCS to access the location server application to the new value that was specified in [Figure 4-3](#). This can be performed via Location Server > Administration > General Properties, as shown in [Figure 4-4](#).

Note that any third-party location clients that have been configured to also use the admin userid to access the location server application via the SOAP/XML API needs to be changed accordingly. You may prefer to define a totally separate userid for each third-party location client that accesses the location appliance, instead of allowing them to use the admin account.

Figure 4-4 Specifying Location Server Application Login Credentials

Cisco Wireless Control System

Monitor ▾ Configure ▾ Location ▾ Administration ▾ Help ▾

Location Server

Administration

- General Properties
- Polling Parameters
- History Parameters
- Advanced Parameters
- Location Parameters
- Notification Parameters
- Active Sessions
- Import Asset Information
- Export Asset Information

Maintenance ▶

Accounts

- Users
- Groups
- Host Access

Status ▶

Logs

Location Server > General Properties > 'AeS_Loc2'

General

Server Name	AeS_Loc2
Version	2.1.34.0
Start Time	6/29/06 4:26 PM
IP Address	10.1.56.21
Contact Name	<input type="text"/>
User Name	<input type="text" value="admin"/>
Password	<input type="password" value="••••"/>
Port	<input type="text" value="8001"/>
HTTPS	<input type="checkbox"/> Enable

190561

Location Appliance Time Synchronization

In order to assure reliable and consistent operation across the network, it is recommended that the WLAN controllers, location appliances and WCS systems within the Cisco UWN maintain synchronized internal clocks. As a general network recommendation, establishing synchronized internal clocks facilitates troubleshooting by making it much easier to correlate log messages between components. Whether viewing independent log files from various components or a combined syslog, having log entries use consistent time stamp references in their message text only serves to make such messages more logical and easier to understand.

This usefulness of consistent timestamps becomes especially clear when multiple location appliances are configured to send asynchronous northbound notifications to a common destination, such as email messages for example. Location appliances configured with the incorrect system time may issue notification messages (as shown in [Figure 4-5](#) bearing incorrect or inconsistent times that may appear confusing to operators at network operations centers (NOCs) or other control points.

Figure 4-5 Email Notification Message Bearing Time Stamp of Location Appliance

```
Date: Sat, 20 May 2006 09:24:01 -0400 (EDT)
From: locserver@st9731.testlab.com
To: wirelessguy@st9731.testlab.com
Subject: TAG ENTERING TEST AREA
X-Mailer: smtpsend
```

Tag 00:0c:cc:5b:ff:44 is in Area Rear Conference Room, Test Lab Annex #2, AP1242 Building, Alpharetta Campus_Group, Alpharetta Campus,

190562

Network Time Protocol (NTP) is the recommended method with which to establish a common clock source and maintain ongoing internal clock synchronization. The Cisco Wireless Location Appliance contains a utility daemon known as *ntpd* that can act as an NTP client to an NTP server located within the enterprise network. A network NTP server provides a common time source reference to all devices, typically using the Coordinated Universal Time (UTC) standard (formerly referred to as Greenwich Mean Time (GMT)).

**Note**

In software releases up to and including Release 4.1, proper time synchronization is recommended. However, with Release 4.2 and beyond, proper time synchronization is mandatory for proper authentication between the location appliance and WLAN controllers.

Complete guidance on configuring and activating the *ntpd* daemon on the location appliance can be found in the following document under the “NTP Configuration and Synchronization for Unified Wireless Network Devices – Set Up NTP on the Location Appliance” section which can be found at the following URL:

http://www.cisco.com/en/US/products/ps6366/products_configuration_example09186a0080811274.shtml#setup-la.

NTP setup information can also be found in the *Cisco 2700 Series Location Appliance Installation and Configuration Guide* found at the following URL:

<http://www.cisco.com/en/US/docs/wireless/location/2700/quick/guide/li31main.html#wp1057105>.

Additional background information and general best practices with regard to NTP in your network may be found in the *Network Time Protocol: Best Practices* document which can be found at the following URL:

http://www.cisco.com/en/US/tech/tk869/tk769/technologies_white_paper09186a0080117070.shtml.

Quiescing the Location Appliance

Although the location appliance is designed to be installed and operated in a continuous fashion, there may be times when it is necessary to power-down the appliance in preparation for extraordinary events such as a physical equipment move or the orderly shutdown of a data center. Simply removing power to the location appliance without undergoing an orderly shutdown may result in any files open at the time becoming corrupted. Although the location appliance’s operating system uses an ext3 journaling file system that minimizes the possibility of file system corruption, it is generally regarded as a best practice to follow the procedure outlined below to initiate an orderly shutdown of all appliance software facilities.

To power-down the location appliance, perform the following steps via either the appliance CLI console or a remote SSH device session.

**Note**

For information on how to connect a CLI console to the location appliance, see “Connecting and Using the CLI Console” section in the *Cisco Wireless Location Appliance: Installation Guide* at the following URL:

http://www.cisco.com/en/US/products/ps6386/products_installation_and_configuration_guide_book09186a00804fa761.html.

Step 1 Manually stop the location server software by issuing the follow command and observing the indicated:

```
# /etc/init.d/locserverd stop
Shutting down locserverd: Request server shutdown now...
Waiting for server...2 secs
Waiting for server...4 secs
.
.
.
.
Waiting for server...60 secs
Server shutdown complete.
#
```

Step 2 Before removing power to the location appliance, issue the following command to properly unmount all file systems, stop all services, and initiate an orderly shutdown of the Linux operating system:

```
# shutdown -h now
```

Issuing this command from the CLI console device in the following output:

```
Shutting down console mouse services: [ OK ]
Stopping sshd:[ OK ]
Stopping xinetd: [ OK ]
Stopping crond: [ OK ]
Saving random seed: [ OK ]
Killing mdmonitor: [ OK ]
Shutting down kernel logger: [ OK ]
Shutting down system logger: [ OK ]
Shutting down interface eth0: [ OK ]
Shutting down loopback interface: [ OK ]
Shutting down audit subsystem[ OK ]
Starting killall: [ OK ]
Sending all processes the TERM signal...
Sending all processes the KILL signal...
Syncing hardware clock to system time
Turning off swap:
Turning off quotas:
Unmounting file systems:
Halting system...
md: stopping all md devices.
flushing ide devices:
Power down.
```

Note that issuing the **shutdown** command from a remote SSH client in your SSH session becoming disconnected. The location appliance still initiates the shutdown procedure, but your SSH session becomes disconnected before the command completes. Therefore, you are not able to view all the command output as you would on a CLI console device. To avoid this lack of visibility, Cisco recommends that a terminal or PC attached to the location appliance console terminal be used to perform this task rather than an SSH session if possible.

Step 3 The final step is to remove power to the location appliance by using the front panel ON/OFF switch to turn the location appliance off. This should be done after the “power down” message is seen on the CLI console (shown in bold above). Note that if using a remote SSH session, you will not see the “power down” message because your session will be disconnected shortly after issuing the shutdown command. In this case, you should wait approximately two minutes for the shutdown command to complete before removing power to the location appliance using the front panel power switch.



CHAPTER 5

Best Practices—Location-Aware WLAN Design Considerations

In the past decade, the design of enterprise-ready wireless LANs has evolved from being centered around the model of maximum coverage with minimum AP count to a model where coverage uniformity and proper cell-to-cell overlap are the predominant concerns. This has been driven by increasing interest in deploying new wireless applications such as wireless voice with its intolerance jitter and high roaming delays. In a similar fashion, deploying location-based applications using a Wi-Fi wireless LAN requires augmenting our traditional approaches, both in the design of “greenfield” location-aware installations as well as the augmentation or retrofit of existing designs.

This chapter describes best practices that should be followed in designing and deploying location-aware wireless LANs and includes the following main sections:

- [Minimum Signal Level Thresholds, page 5-2](#)
- [Access Point Placement, page 5-5](#)
- [Access Point Separation, page 5-12](#)
- [Determining Location Readiness, page 5-18](#)
- [Location, Voice and Data Coexistence, page 5-20](#)
- [Avoiding Location Display Jitter, page 5-31](#)
- [Multiple Location Appliance Designs, page 5-32](#)
- [Antenna Considerations, page 5-44](#)
- [Calibration, page 5-48](#)
- [Inspecting Location Quality, page 5-64](#)
- [Using Test Points to Verify Accuracy, page 5-68](#)

Minimum Signal Level Thresholds

For mobile devices to be tracked properly, a minimum of three access points (with four or more preferred for better accuracy and precision) should be detecting and reporting the received signal strength (RSSI) of any client station, asset tag, or rogue device being tracked. It is preferred that this detected signal strength level be -75dBm or better.

**Note**

As of WLAN controller software Release 4.1.185.0, each tracked entity (WLAN client, RFID tag, rogue access point, or rogue client) is detected by up to sixteen registered access points at any time on each WLAN controller. This helps to improve the tracking of devices in motion across many access point coverage cells by assuring that the latest device RSSI is properly reflected.

When performing a site survey of an area where clients or tags are tracked, the RSSI of representative devices should be verified to ensure compliance with the minimum number of recommended access points and the recommended detected signal strength. This should be performed via one of two techniques:

- Viewing detected RSSI for the client or asset tag using the **show client detail** *<mac address>* or **show rfid detail** *<mac address>* controller CLI command, as shown in [Figure 5-1](#).
- Viewing detected RSSI for the client or asset tag using the location floor map GUI, as described in [Figure 5-2](#) and [Figure 5-3](#).

Figure 5-1 Checking Client RSSI at the WLAN Controller

```
(Cisco Controller) >show client detail 00:07:50:d5:e4:77
Client MAC Address..... 00:07:50:d5:e4:77
Client Username ..... N/A
AP MAC Address..... 00:14:1b:59:41:f0
Client State..... Associated
Wireless LAN Id..... 1
BSSID..... 00:14:1b:59:41:f0
Channel..... 1
IP Address..... 10.1.67.252
Association Id..... 1
Authentication Algorithm..... Open System
Reason Code..... 0
Status Code..... 0
Session Timeout..... 0
Client CCX version..... 2
Mirroring..... Disabled
QoS Level..... Silver
Diff Serv Code Point (DSCP)..... disabled
802.1P Priority Tag..... disabled
WMM Support..... Disabled
Mobility State..... Local
Mobility Move Count..... 0
Security Policy Completed..... Yes
Policy Manager State..... RUN
Policy Manager Rule Created..... Yes
NPU Fast Fast Notified..... Yes
Policy Type..... N/A
Management Frame Protection..... No
EAP Type..... Unknown
Interface..... vlan64
VLAN..... 64
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..... 2816
  Number of Bytes Sent..... 872
  Number of Packets Received..... 19
  Number of Packets Sent..... 6
  Number of Policy Errors..... 0
  Radio Signal Strength Indicator..... -45 dBm
  Signal to Noise Ratio..... 55 dB
Nearby AP Statistics:
  TxExcessiveRetries: 0
  TxRetries: 0
  RtsSuccessCnt: 0
  RtsFailCnt: 0
  TxFiltered: 0
  TxRateProfile: [0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0]
AP1242#1(slot 0) .....
antenna0: 82 seconds ago -63 dBm..... antenna1: 82 seconds ago -68 dBm
AP1242#2(slot 0) .....
antenna0: 82 seconds ago -54 dBm..... antenna1: 82 seconds ago -59 dBm
AP1242#3(slot 0) .....
antenna0: 82 seconds ago -45 dBm..... antenna1: 82 seconds ago -41 dBm
AP1242#4(slot 0) .....
antenna0: 82 seconds ago -65 dBm..... antenna1: 82 seconds ago -68 dBm
AP1242#5(slot 0) .....
antenna0: 82 seconds ago -64 dBm..... antenna1: 82 seconds ago -66 dBm
```

In either case, these techniques should be used with representative test clients or asset tags in the area where localization is desired. When performing this check, it is important to ensure that all access points and antennas are installed and representative of the final configuration. The maximum transmit power level supported as well as the probing behavior of the test client should be as close as possible to that of the production clients you wish to track. Figure 5-1 indicates that the output of the CLI command displaying the signal strength of the client as detected by all of the access points detecting the client, registered to the same controller. In situations where the detecting access point registrations are distributed among two or more controllers, more than one CLI session is required. From the information provided within the red rectangular area in Figure 5-1, it can clearly be seen whether or not the client in

question is being detected by three or more access points at the recommended signal strength level or better. In a similar fashion to that shown for WLAN clients in Figure 5-1, the CLI command **show rfid detail <mac address>** can be used to display detected RSSI information for an asset tag.

This same information can be obtained graphically via the location map GUI by clicking on either a WLAN client icon (blue rectangle) or asset tag icon (yellow tag), enabling the location debug checkbox and then enlarging the miniature location map as shown in Figure 5-2 and Figure 5-3.

Figure 5-2 Enabling Location Debug

The screenshot displays the 'Tags > Tag Asset' configuration page. At the top right, there is a dropdown menu set to '-- Select a command --' and a 'GO' button. The page is divided into several sections:

- Tag Properties:** A table with the following data:

Vendor	Aeroscout
Controller	10.1.56.18
Battery Life	Normal
- Location:** A table with the following data:


Floor	Alpharetta Campus_Group>AP1242 Building>Test Lab Annex #2
Last located at	May 2, 2006 10:03:19 PM
On Location Server	AeS_Loc2
- Asset Info:** A form with input fields for Name, Group (AeroScout RFID), and Category. Below these is a 'Location Debug' checkbox which is checked and labeled 'Enabled*'. An 'Update' button is located below the checkbox. A note below the form states: '* This will show AP RSSI Information on the Map.'
- Statistics:** A table with the following data:

Bytes received	101040
Packets received	3368
- Location Notifications:** A table with the following data:

Absence	0
Containment	0
Distance	0
All	0

At the bottom left of the location map area, there is an 'Enlarge' link. On the right side of the page, the number '190564' is printed vertically.

Figure 5-3 *Displaying Detected RSSI via the GUI*



Access Point Placement

Proper placement of access points is one of several best practices that should be adhered to in order to unleash the full performance potential of the location-aware Cisco Unified Wireless Network. In many existing office wireless LANs, access points are distributed mainly throughout interior spaces, providing service to the surrounding work areas. These access point locations have been selected traditionally on the basis of coverage, WLAN bandwidth, channel reuse, cell-to-cell overlap, security, aesthetics, and deployment feasibility. In a location-aware WLAN design, the requirements of underlying data and voice applications should be combined with the requirements for good location fidelity. Depending on the particular site, the requirements of the location-aware Cisco UWN are flexible enough such that the addition of location tracking to voice installations already designed in accordance with Cisco best practices, for example, may not require extensive reworking. Rather, infrastructure already deployed in accordance with accepted voice best practices can often be augmented such that location tracking best practice requirements are met as well, (such as perimeter and corner access point placement, for example) depending on the characteristics of the areas involved.

In a location-ready design, it is important to ensure that access points are not solely clustered in the interior and toward the center of floors. Rather, perimeter access points should complement access points located within floor interior areas. In addition, access points should be placed in each of the four corners of the floor, and at any other corners that are encountered along the floor perimeter. These perimeter access points play a vital role in ensuring good location fidelity within the areas they encircle, and in some cases may participate in the provisioning of general voice or data coverage as well.

If using chokepoint location, verify that all areas planned for chokepoint trigger installation are clearly within the range of your access points. In addition to ensuring that messages transmitted by asset tags located within chokepoint areas are properly received by the system, proper planning can help assure


that asset tags can be tracked using RF Fingerprinting as they approach and exit chokepoints. The ability to track asset tags using RF Fingerprinting complements the system's ability to locate tagged assets within chokepoint areas using highly granular chokepoint location techniques.

The access points that form the perimeter and corners of the floor can be thought of as outlining the *convex hull* or set of possible device locations where the best potential for high accuracy and precision exists. By definition, the convex hull of a set S of points, denoted $hull(S)$, can be regarded as the smallest polygon P for which each point of S is located either on the boundary or within the interior of P .

Figure 5-4 illustrates the concept of a convex hull. In Figure 5-4, assume the set of access point locations is denoted by the black dots, which we refer to as *set S*. The convex hull of set S , or $Hull(S)$, is figuratively represented as an elastic band (shown by the blue line) that is stretched and allowed to snap over the outermost members of the set (which in this case represents perimeter and corner access points).

The interior area encompassed by this band (depicted in green) can be considered as possessing high potential for good location accuracy. As tracked devices stray into the area outside the convex hull (outside the green area in Figure 5-4), accuracy can begin to deteriorate. Although it may vary given the number of access points deployed and their inter-access point spacing, generally speaking, the rate of this accuracy degradation has been seen to be almost linear as the tracked device moves further and further outside the convex hull. For example, a device that experiences less than or equal to 10m/90% accuracy within the convex hull may deteriorate to 18m/90% by the time the device moves to a point 20 feet outside it.

Figure 5-4 *The Convex Hull of a Set of Points*

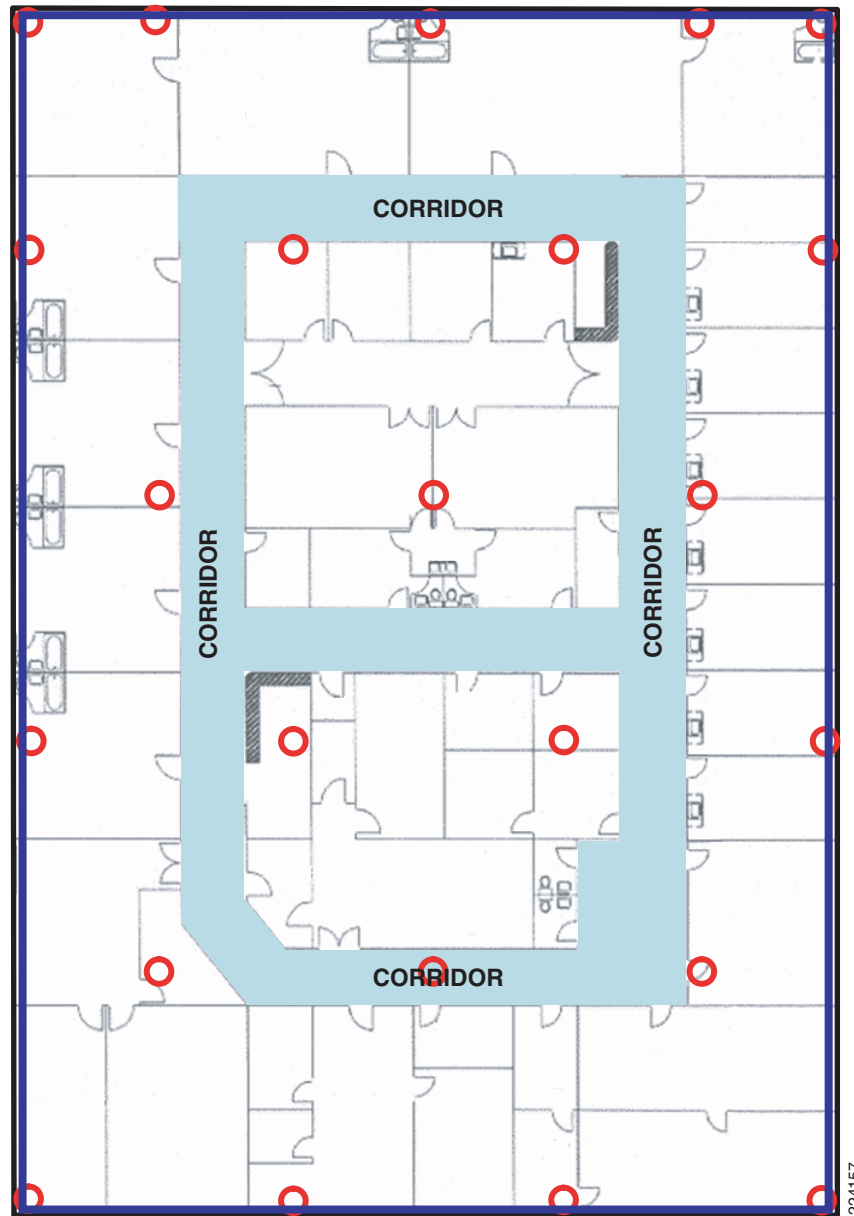


In order to assure proper convex hull establishment around the set of location data points possessing high potential for good accuracy, access points should be placed in each corner of the floor, as well as along the floor perimeter between corners. Inter-access point separation along the perimeter should be in accordance with the general access point separation guidelines (described in a subsequent section). The designer may reduce this spacing if necessary, in order for these access points to participate in the provisioning of voice or data service to the floor.

Figure 5-5 provides an illustration where these concepts are applied to a floor with a type of floor plan found in many enterprises (that of rooms or offices contained by and surrounding an interior corridor). In this case, the area in which we desire to locate tracked assets is the entire floor. In Figure 5-5, note that the access points located towards the center of the floor are complemented by those that have been

placed along the perimeter. As is the case in most proper location-aware designs, the set of location data points possessing the highest potential for good location accuracy is contained within the convex hull, which in Figure 5-5 is represented by the blue rectangle and encompasses the entire floor.


Figure 5-5 Proper Access Point Perimeter Placement



In some cases, customer preferences or deployment restrictions may factor into the access point placement decision, and the placement of access points at the floor perimeter may be restricted in one way or another. While this still may result in acceptable placement from the perspective of providing basic RF coverage, because there may be significant areas where asset tracking is required outside the access point perimeter (and thus outside the convex hull), such placement may lead to reduced location fidelity in those areas.

Figure 5-6 illustrates an example of a less-than-desirable situation where the placement of access points has been restricted to hallway corridors and administrative/storage facilities located within the areas encircled by the corridors. For aesthetic reasons, facilities management has decided that access points will not be placed within any of the executive offices or conference rooms located between the hallway corridors and the physical perimeter. Because of these restrictions, our convex hull now lies at the outside edge of the corridor (indicated by the blue rectangle) and not at the true physical perimeter of the floor.

Figure 5-6 Artificially Constrained Access Point Perimeter



224158


Given what we know about the distribution of location errors when operating outside the convex hull, it is logical to expect that location accuracy will not be as good in the offices and rooms located there. These areas of potentially lower accuracy are highlighted in red in Figure 5-6.

With our recommendation of establishing the convex hull at the true floor physical perimeter notwithstanding, in practice the difference in location error rate between points located within the convex hull and outside it may be tolerable in some situations. These might include situations where such areas extend beyond the office perimeter for only a short distance (for example, small 10x10 foot rooms lining the walls of a corridor). For example, looking at the areas highlighted in red in [Figure 5-6](#), the potential increase in location error would be less in the smaller offices located at the right side of the floor plan than in any other affected area. Depending on magnitude, the effect of operation outside the convex hull will likely be the least. In contrast, the areas at the bottom of the floor plan, with larger offices and multiple wall partitions, would be potentially effected to a significantly higher degree.

In cases where access point placement in perimeter offices and conference rooms is restricted due to aesthetic concerns, a potential compromise may be possible using a very low profile antenna (such as the Cisco AIR-ANT5959 or Cisco AIR-5145V-R) along with access point mounting in a plenum-rated enclosure (where permitted by local codes). This would offer the ability to mount access points at the proper perimeter and corner locations (thereby avoiding the quandary described in [Figure 5-6](#)), but with minimal visible footprint to the casual observer.

As mentioned earlier, the floor plans shown in [Figure 5-5](#) and [Figure 5-6](#) are commonplace, but by no means exclusive. For example, some modern building designs may possess hallway corridors that are located directly alongside the actual floor and building perimeter, typically allowing a panoramic view of campus environs as visitors move about between offices and conference rooms. In this case, all offices and conference facilities are located within the area between the corridors and the center of the floor. [Figure 5-7](#) provides an illustration of such a floor plan. Note that with this floor layout, placement along the outer edge of the hallway corridor places the access points along the actual physical perimeter, by default.

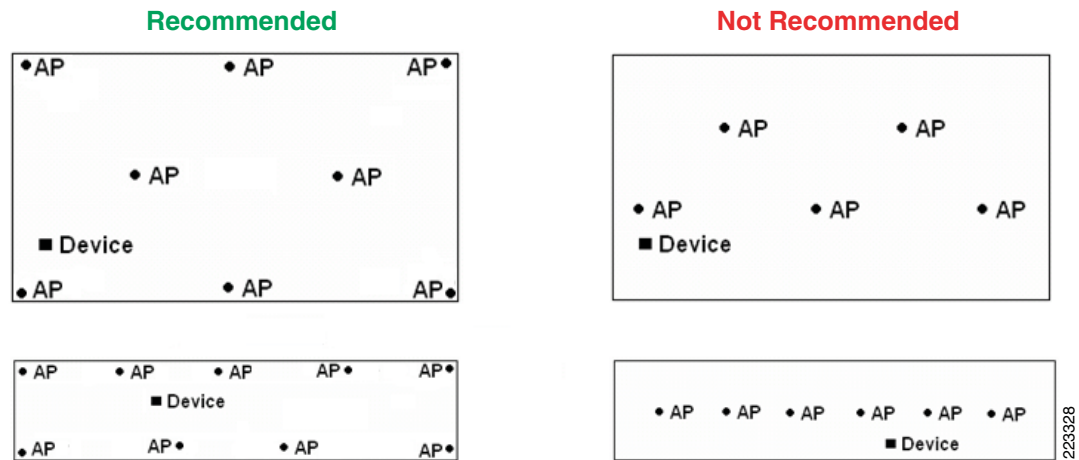
Figure 5-7 Perimeter Corridor Floor Plan



224159

Figure 5-8 provides simple illustrations summarizing the access point placement concepts discussed in this section so far. Note that designs that make use of only clustered or straight-line access point placement should be augmented or redesigned in favor of those that combine center access-point placement with perimeter and corner placement.

Figure 5-8 Basic Example of Location-Aware Access Point Deployment



If possible, mount antennas such that they have an unencumbered 360° view of all areas around them, without being blocked at close range by large objects. For example, if possible, avoid placing access point antennas directly against large objects such as steel columns, as illustrated in Figure 5-9. One option is to mount the access point along with its antennas to a ceiling location (provided that this allows an acceptable mounting height). Another option is to use short, low loss cable extension to allow separation between antennas and such obstructions.

Figure 5-9 Access Point Mounted Directly to Steel Column



Additional discussion of proper access point placement can be found in *Cisco Wireless Location Appliance: Deployment Guide* at the following URL:
http://www.cisco.com/en/US/products/ps6386/prod_technical_reference09186a008059ce31.html.


Access Point Separation

The distance between deployed access points can impact location performance, as well as the performance of co-resident voice and data applications. From a location perspective, while location tracking inter-access point spacing requirements tend to be relatively flexible and supportive of the coverage needs of underlying applications, very small or very large inter-access point separation distances are usually best avoided.

An excessive inter-access point distance¹ can detract from good location accuracy by not providing sufficient signal strength differentiation at extended distance. Insufficient inter-access point distance can expose the system to short range antenna pattern anomalies, which may also be non-conducive to good location accuracy. From the perspective of co-resident voice and data applications, the inter-access point distance is one of the key factors determining whether required minimum signal level thresholds, data rate thresholds, signal to noise ratios (SNR), and required coverage overlap will be met. From a location accuracy perspective, the range of acceptable inter-access point distance tends to be rather broad, and can provide excellent location accuracy while accommodating the needs of most co-resident voice and data applications.

The techniques incorporated in the location-aware Cisco UWN to localize tracked devices operate most effectively when RSSI and distance are seen to possess a clearly *monotonic* relationship. To better understand what is meant by this, we examine a simulated plot of a tracked device's detected RSSI as the distance between it and a detecting access point is increased (see Figure 5-10). While the relationship between RSSI and distance varies depending on different combinations of antenna, antenna height and environmental characteristics, the graph shown in Figure 5-10 for an access point mounted at approximately twelve feet elevation can be used to better understand the concepts discussed here.

Figure 5-10 An Example of the Relationship Between RSSI and Distance



1. As discussed in a later section of this design guide, excessive antenna heights can also contribute to diminished accuracy.


In [Figure 5-10](#), we see that beginning at some point *a* fairly near the access point, and ranging to another point *c* further in distance, the two variables exhibit a strict monotonically decreasing relationship (as distance between the tracked device and the access point increases, the RSSI at which the access point detects the device is shown to decrease). Between point *a* and another point *b*, the amount of change in RSSI (dBm) that occurs per-unit change in distance (feet) is highly consistent, approximately -5 dBm per 20-foot change in distance. This results in the slope of the graph between points *a* and *b* being fairly steep. As the distance continues to increase beyond point *b*, the slope of our graph begins to diminish and the level of RSSI differentiation decreases, providing increasingly less differentiation in received signal strength per-unit change in distance. Note that the slope of the graph between points *b* and *c* is not nearly as steep as it is between points *a* and *b*. As distance begins to significantly exceed point *b* in this example, the slope of the graph will diminish even further. This greatly reduced slope and steepness results in a decreased level of differentiation in signal level with increasing distance. When this occurs at extended distances, it becomes more difficult to accurately predict changes in distance based on detected changes in RSSI (lateration).

The risk of this lack of RSSI differentiation having a significant impact on location accuracy can be reduced if steps are taken to avoid areas of the RSSI versus distance curve where this phenomena is known to exist most prominently. In general, for access points deployed indoors at antenna heights of 20 feet or less, this can be achieved if the range of any point on the floor to at least three detecting access points on that floor (one in each of at least three of the four quadrants surrounding it) is maintained within approximately 70 feet in an indoor environment. This is a general recommendation that is intended to assist designers in avoiding situations where excessive inter-access point distance may be a contributing factor to location inaccuracy. As shown in [Figure 5-10](#), diminished RSSI differentiation with increasing distance is a gradually increasing phenomenon, therefore, a degree of flexibility is implied in this recommendation.

In practice, in addition to being conducive to good location accuracy, this recommendation applies well to deployments where location tracking is deployed in conjunction with other WLAN applications (such as voice and high speed data) in accordance with current recommended best practices. This is especially true for environments where the expected path loss exponent is 3.5 (walled office environment) or higher, as the required inter-access point spacing tends to generally fall within this range. In addition to the potential effects of a lack of RSSI differentiation at distance extremes, inter-access point distances significantly greater than 70 to 80 feet can make it more challenging to satisfy the best practice signal strength and overlap requirements of VoWLAN devices such as the Cisco 7921G and the Vocera Communication Badge in environments with high path loss.

At ranges closer than point *a* in our example, propagation anomalies that are due to the elevation pattern of the chosen antenna, the antenna's installation height, and the current physical location of the tracked device can potentially combine to degrade monotonicity. As a result, RSSI cannot be depended on as a reliable predictor of distance in this part of the curve, since it may be possible that more than one equally likely value for distance exists at a particular detected RSSI level. [Figure 5-11](#) illustrates this case, depicting how a tracked device's RSSI reading of -40dBm can be associated with three different distances (5, 7, and 12 feet) from the access point antenna when operating in this close-range non-monotonic region of the RSSI versus distance graph. This behavior is typically the result of a variation in an overhead antenna's propagation pattern as a device approaches it begins to venture into the area almost directly beneath it. Obviously, these effects vary depending on the propagation pattern of the specific antennas used and their installation height above the area where tracked devices is located. However, the lesson to be learned from this is that although increased access point density can often be conducive to better location accuracy, the effect is not without its limits.

Figure 5-11 Example of Close-Range Non-Monotonicity



Clearly, such RSSI ambiguity can be confusing, especially when attempting to use RSSI to accurately laterate distance. Such ambiguous behavior is generally not conducive to good location fidelity. In tests conducted with access points at an installed height of 10 feet in with 2.2dBi omni-directional antennas in an environment with a path loss exponent of 3.4, this behavior could sporadically be observed out to a distance of almost 14 feet. In the specific case of this example, it would be best to maintain the inter-access point spacing above 28 feet (in other words, twice the distance at which such behavior would be expected) in order to reduce the potential of this phenomena occurring.

In some application designs, it may be desirable to deploy multiple access points on non-overlapping channels in order to potentially increase the amount of RF bandwidth available to users¹ (“collocated non-overlapping access points”). This approach is often seen in classrooms and conference halls where there may be a large number of mobile users. If location tracking of WLAN clients and other devices is desirable in situations where some rooms may possess several collocated access points, it is suggested that the co-located access points not be deployed within very close proximity (i.e. a few feet) of each other. Rather, every attempt should be made to obtain as much separation as possible between these co-located access points, so as to avoid any of the close-range effects that can be detrimental to good location fidelity. One way to accomplish this for co-located access points in a lecture hall, for example, would be to place the access points on different walls and perhaps the ceiling as well, with appropriate inter-access point spacing.

In general then, most indoor location tracking deployments with access point antennas installed at heights of between ten and twenty feet can be well served with an inter-access point spacing of between 40 and 70 feet, especially when combined with the signal threshold and access point placement recommendations suggested in the preceding sections of this document. In some cases however, inter-access point spacing below 40 feet may be necessary to satisfy the requirements of some applications for high signal strength thresholds, especially in environments where high path loss is present. An example of this might be a voice application deployed in such an environment (for example, a path loss exponent of 4.0 where a high degree of environmental clutter is present). Best practices for Cisco 7921G VoWLAN deployments would suggest a minimum signal level of -67dBm, 20% inter-cell overlap and signal to noise ratio of 25 dB for 802.11g in this type of situation. Applying these requirements mathematically, we calculate an estimated cell size of 24 feet and an inter-access point spacing of 33 feet. In this case, in order to deploy our voice application in accordance with recommended

1. Note that any realized increase in bandwidth from co-located access points is subject to limitations due to co-channel interference from other access points on those same channels.

best practices, the inter-access point spacing should be reduced below the general guideline of 40 feet. Note that good location accuracy is achievable at inter-access point ranges below 40 feet, provided that the access point spacing is not decreased so much that the negative effects of close range non-monotonicity come into play. Generally, this should not be an issue if the inter-access point distances are above 25 to 28 feet when using low gain, omni-directional antennas mounted at an installation height of approximately 10 feet in an indoor environment.

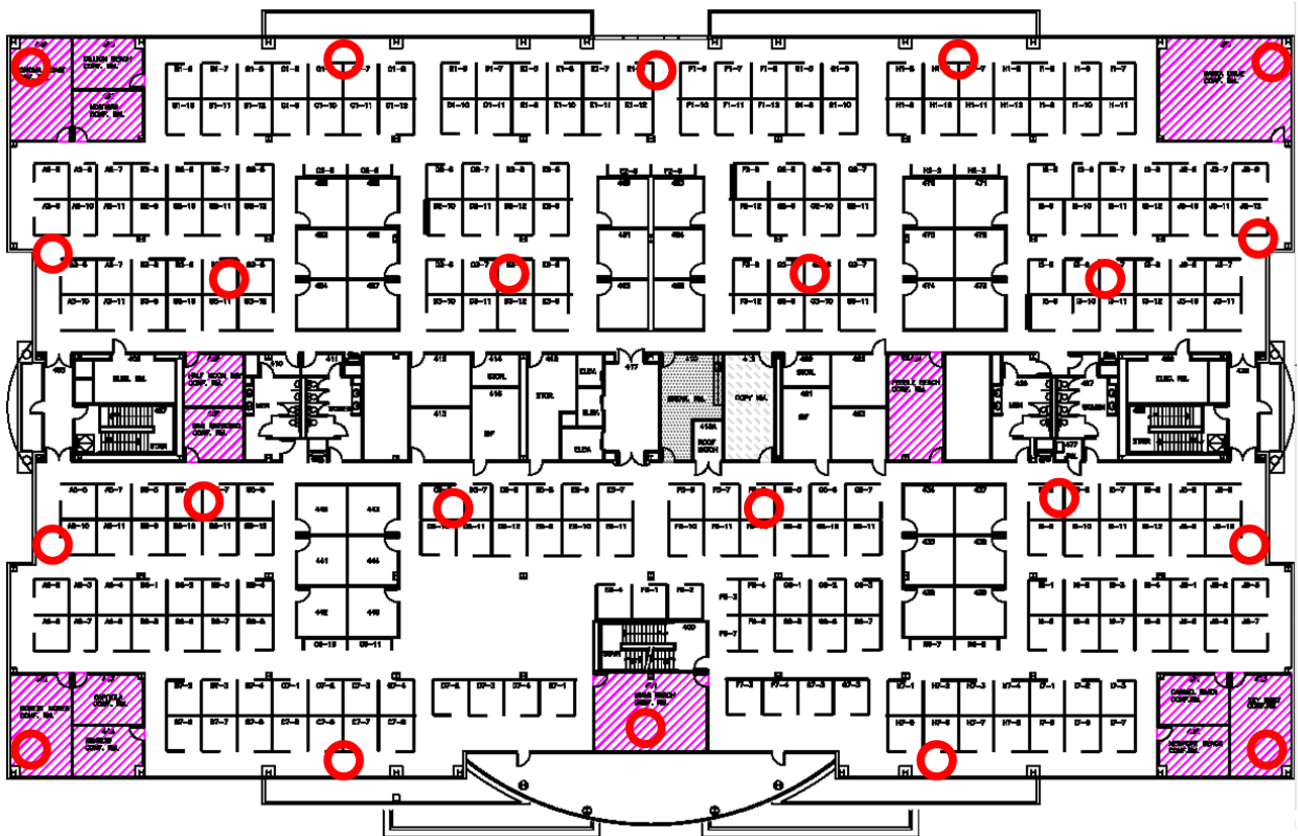
Figure 5-12 illustrates an example of access point placement and inter-access point spacing, offering a foundation for a location-aware design. The environment in Figure 5-12 consists of drywall offices and cubicle office spaces with a total space of approximately 275 feet by 159 feet. Taking into consideration the location tracking requirement for illustrative purposes *only*, our inter-access point linear-spacing recommendations of 40 to 70 feet suggests approximately 22 location-aware access points as an initial estimate. Incorporating the placement strategies made in preceding sections, interior, perimeter, and corner access points are placed to facilitate multi-lateration and establish a clearly delineated convex hull around the floor.



Note

In an actual installation involving WLAN applications deployed in conjunction with location tracking, interior access point design should be conducted prior to instituting design modifications in support of location tracking modifications to ensure that best practice recommendations for signal strength, overlap and signal to noise ratio requirements of data and voice applications are met.

Figure 5-12 An Example of Location Aware Access Point Placement



WCS includes a planning tool that allows designers to model “what-if” design scenarios. The WCS Planning Tool is accessible via the Monitor > Maps > *floormapname* > Planning Mode dropdown menu selection. This is a predictive modeling tool that is used on a per-floor basis to provide initial guidance on access point placement, as well as an interactive representation of predicted access point signal strength and data rate information. It can be safely used without impacting any actual deployment of access points that may already be in service. The WCS Map Editor is accessible from the top line hyperlink bar of the planning tool, and can be used to add wall attenuation information to floor maps. Wall information added via the Map Editor does not affect access point placement or location designs, however, it will be used by the planning tool when displaying predicted RF coverage maps for planned access points.

The planning tool operates purely on a hypothetical basis without the need to connect or deploy any access points or controllers. Since it is WCS feature, a WCS server must be installed somewhere in network before the planning tool can be used. If there are any existing access points that have been deployed and defined to WCS already, the planning tool allows for the configuration of those access points to be copied into the planning virtual environment, allowing you to safely model with a virtual copy of your production environment.

Before using the planning tool for RF coverage planning, ensure that an appropriate path loss model has been assigned to the floor upon which you wish to conduct your planning. WCS will use the coverage reference path losses and path loss exponents when it plots the predicted coverage heatmaps from each access point in the planning tool. Seasoned WLAN veteran designers have the option of using the planning tool in a manual mode to place access points on floor maps as they see fit and adjust several criteria in order to see their effect (such as transmit power, antenna type, and so on). Alternatively, the WCS planning tool also allows automated access point placement based on the type of deployment model desired. Those users and designers desiring that the system make an initial design suggestion can use the planning tool in an automated mode, thereby specifying the type of design they wish and allowing the planning tool to examine their requirements and make qualified suggestions. For designers wishing to combine voice and data designs meeting Cisco VoWLAN best practices with location tracking, it is recommended that the planning tool be first used to model voice and data designs separately from location tracking requirements. Once a satisfactory voice and data design has been created, any modifications necessary to provide for good location fidelity can then be manually incorporated.

The planning tool assumes a transmit power of +18dBm for 802.11bg and +15dBm for 802.11a, along with an antenna azimuth position of 180°, elevation height of ten feet and elevation angle of 0°. Transmit power, access point type, antenna type, and azimuth position can be changed individually for each access point. In addition, planning tool users can specify a several additional criteria to further fine tune data and voice designs.

**Note**

For complete information about planning tool options (such as data/coverage, voice, location, demand and override) consult the chapter *Using Planning Mode to Calculate Access Point Requirements* found in the document entitled *Cisco Wireless Control System Configuration Guide, Release 4.1* at the following URL:

<http://www.cisco.com/en/US/docs/wireless/wcs/4.1/configuration/guide/wcsmaps.html#wp1104248>.

Selecting the location planning option results in the planning mode access points being placed along the perimeter and in the corners of a floor, in addition to the interior of the floor as necessary. At least four access points are assumed to be present in every location design, and access points are placed using a spacing of up to 70 feet. Note that when using the location planning option, the resulting design may meet best practice recommendations for voice and data, although the signal strength and overlap requirements of co-resident applications are not explicitly taken into account. Therefore, in designs where location tracking is intended to co-reside with voice and high speed data, it is recommended that


these application designs be addressed first, according to Cisco-recommended best practices. Once a design satisfying application needs has been completed, the design can then be modified or augmented as necessary to meet location tracking requirements.

Prior to software Release 4.1, the automated placement capabilities of the planning tool were limited to floors and buildings whose shapes were simple polygons, such as squares and rectangles. A new capability added in this release allows the planning tool to accommodate irregularly shaped floor areas. To accomplish this, an irregular coverage perimeter is drawn using the WCS Map Editor and when saved, becomes available for use in the planning tool. When designing new floor layouts using the planning tool, the user is allowed to choose between using the traditional closed polygon or the newly created irregular shape. Further information on this newly introduced capability can be found in the chapter *Using the Map Editor to Draw Polygon Areas* in the document entitled *Cisco Wireless Control System Configuration Guide, Release 4.1* at the following URL:

<http://www.cisco.com/en/US/docs/wireless/wcs/4.1/configuration/guide/wcsmaps.html#wp1104253>.


With this enhancement, the planning tool becomes more useful to designers when working with common buildings having irregular shapes, such as a building with an open courtyard as shown in [Figure 5-13](#). In [Figure 5-13](#), we see a location design performed using the automatic planning tool mode. Note the red outlined perimeter of the building, which was added to the floor image using the Map Editor and is now eligible for use within the planning tool.

Figure 5-13 Using WCS Planning Tool with an Irregularly Shaped Floor Plan



More complex designs containing totally enclosed interior voids (for example, a building with a fully enclosed interior atrium as shown in [Figure 5-14](#), with the perimeter of the building shown by a red outline) may not lend themselves well to automatic access point placement. The planning tool does not currently allow the exclusion of zones into which access point placement should not occur. Note in [Figure 5-14](#) the placement of access points 2, 4, 9 and 24 in the atrium area (indicated by the blue outline). The placement of these access points in this area is incorrect, since the floor map is for the building's third floor. This should be corrected by manual intervention and moving the access points into correct locations or eliminating them entirely if not necessary.

Figure 5-14 Example of Floor Plan with Fully Enclosed Interior Atrium



2230383

Determining Location Readiness

The Inspect Location Readiness feature (Monitor > Maps > *floormapname* > Inspect Location Readiness) allows the network designer to perform a quick predictive check of the location performance for a floor before time is invested in pulling cable, deploying equipment, and performing calibrations. Inspect Location Readiness takes into consideration the placement of each access point along with the inter-access point spacing indicated on floor maps to predict whether estimated location tracking accuracy will be within 10 meters in 90 percent of all cases. The output of the location readiness inspection is a “go / no-go” graphical representation of the areas that are predicted to be likely candidates for producing this level of accuracy, as well as those that are not.

- Note that unlike the planning tool described earlier, the location readiness tool assumes that access points and controllers are known to WCS and have been defined on the WCS floor maps using Monitor > Maps > Position APs. While it is not necessary to actually install access points and antennas on walls and ceilings in order to conduct a location readiness assessment, you must add any applicable controllers to WCS along with their registered access points, and place the icons representing the access points on the appropriate floor maps. In order to do this initially, all such controllers and access points must be physically present, powered on and online to the network. These access points and controllers need not be deployed and installed, they can be on a floor,

tabletop or other temporary location, as long as the access points and controllers are capable of communicating with WCS. The Cisco Wireless Location Appliance need not be present in order to conduct a location readiness assessment.

- Once the access points that you wish to place on floor maps have been added to the WCS database, subsequent location readiness assessments can be conducted using these same access points, even if they are not reachable from WCS at that time. Because the location readiness inspection is based on access point placement and the inter-access point distances shown on the floor maps, accurate map placement of access points is very highly recommended. The location readiness tool is used to only assess the preparedness of the design to perform RF Fingerprinting-based location tracking. It does not validate any aspect of the design to perform chokepoint location, especially with regard to the definition or positioning of chokepoint triggers. After access point placement has been performed, select the floor map that you wish to verify the location readiness of and then choose Inspect Location Readiness from the upper right-hand dropdown command menu.

A point is defined as being “location-ready” if the following are all determined to be true:

- At least four access points are deployed on the floor
- At least one access point is found to be resident in each quadrant surrounding the point-in-question
- At least one access point residing in each of at least three of the surrounding quadrants is located within 70 feet of the point-in-question.

Figure 5-15 illustrates our three location readiness rules, where the green and yellow circles represent access point locations and the point-in-question is represented by a red dot.

Figure 5-15 Definition of a “Location-Ready” Point




Figure 5-16 shows an example of a floor deployment where not all areas have passed three-point location readiness assessment described earlier for 10m/90% accuracy. Although there are green areas toward the center of the figure, notice that red areas abound as you get beyond the perimeter access points representing the convex hull. By establishing a solid understanding of the requirements that define location readiness, the information contained in Figure 5-15 can be used to help determine how access points may be required to be relocated (or additional access points introduced) to improve performance. For example, if 10m/90% or better location accuracy is required within the red areas, additional access points could be introduced to establish a more clearly delineated floor perimeter, including the placement of access points in the corners of the floor and re-checking inter-access point distances. By implementing these types of modifications, the ability of the Cisco UWN to resolve the location of tracked devices in these highlighted areas is likely to be significantly enhanced.

Figure 5-16 Example of Location Readiness Tool Usage



Once again, keep in mind that location readiness inspection is a distance-based *predictive* tool. As is the case with most predictive tools, it can be expected that some degree of variance will occur between predicted and actual results. Cisco recommends that the location readiness tool be used in conjunction with other best-practice techniques outlined in this document, including the location quality inspection.

Location, Voice and Data Coexistence

The location-aware Cisco Unified Wireless Network is a multi-purpose wireless platform that allows enterprises to bring consistency and efficiency to their business processes, providing increased overall effectiveness. A key advantage of the location-aware Cisco UWN is the integration and the cost advantage that stems from its ability to perform high quality location tracking of clients, asset tags and rogue devices with only reasonable additional investment required beyond that necessary to support other enterprise wireless applications, such as VoWLAN and high speed data.



Note

This section describes the pertinent characteristics of voice and data designs only as they relate to co-existence with the location tracking capabilities of the Cisco UWN. For a more comprehensive examination of Cisco Unified Wireless Network VoWLAN solution design and Cisco recommended best practices, refer to the *Voice Over Wireless LAN 4.1 Design Guide* which can be found at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>.

**Note**

For a more comprehensive examination of the Cisco Unified Wireless Network data solution design and best practices, refer to *Enterprise Mobility 4.1 Design Guide* at the following URL:
<http://www.cisco.com/go/srnd>.

When architecting VoWLAN and high speed data designs and determining subsequent access point placement, the primary concerns of the designer should include:

- Minimum desired cell signal level threshold—For example, when designing VoWLAN solutions that involve the Cisco 7921G VoWLAN handset, current VoWLAN best practices suggest a minimum planned signal level threshold of -67dBm. Other voice devices may have differing requirements (such as the Vocera Communications badge, which requires a signal level threshold of -65dBm). Requirements for data devices will depend on the transmission rate that they are required to operate at. Lower speed devices (such as handheld bar code or RFID computers that operate at data rates up to 11 Mbps) typically do not have very demanding minimum signal requirements, often times in the range of -73 to -76 dBm. Data devices used to pass streaming multimedia and other bandwidth-intensive applications will typically require higher data transmission rates and consequently, higher minimum signal levels.
- Signal to Noise Ratio (SNR) —This is the ratio of the signal strength at the receiver to the noise floor, and is measured in dB. Since both components of the ratio are specified in dBm, the SNR can be calculated by simply subtracting the noise value from the signal strength value. The minimum required SNR for a receiver to operate properly varies depending on construction of the receiver, as well as the bit rate or modulation it is expected to operate at. A typical example is shown below:

Transmission Rate (Mbps)	1	2	5.5	11	6	9	12	18	24	36	48	54
Signal to Noise Ratio (dB)	4	6	8	10	4	5	7	9	12	16	20	21


Ensuring the existence of sufficient SNR is very important when designing for robust and reliable wireless application support. This is especially so in wireless voice applications, where it is necessary to ensure that a high percentage of packets are successfully decoded in each cell and jitter is kept to a minimum. For example, with a Cisco 7921G VoWLAN handset the recommended SNR to ensure a good user VoWLAN experience is 25 dB. Keep in mind that if the signal to noise ratio is insufficient due to a high noise floor, proper operation of the wireless device may be difficult to achieve in spite of high overall received signal levels. SNR and minimum received signal levels should be considered together in order to assure that a new deployment has met design standards and is ready for production pilot testing.

- Data rate—Data bit rates are enabled or disabled via the wireless infrastructure, with minimum signal level thresholds and the signal to noise ratio determining which of the enabled bit rates will actually be usable. For example, with the Cisco 7921G, the combination of a -67dBm minimum signal level and a 25dB signal to noise ratio generally makes the use of 24 Mbps or greater data rates possible.
- Cell-to-cell overlap—In a very simple sense, we can think of each of our access points as residing at the center of an RF “cell” with a spherical boundary of RF coverage around them. Our primary interest is in the coverage boundary associated with our desired minimal signal threshold. In order to provide consistent coverage and availability across our floor, each of our cells should join with each adjacent cell at a coverage boundary that is greater than our desired minimal signal threshold. How much greater? That is determined by the amount of cell-to-cell overlap we wish to implement in our design, which in conjunction with the other parameters we have described, will dictate the potential packet loss experienced by VoWLAN devices before a roam event occurs.

The application of cell-to-cell overlap is intended to increase the probability that VoWLAN clients will quickly detect and roam to an adjacent cell without enduring an excessive degree of rate shifting and re-transmission as the device approaches the cell boundary. Excessive rate shifting and packet re-transmission is especially counter-productive for VoWLAN devices, as such behavior typically results in packet loss which usually translates into jitter. Since jitter is well established to be detrimental to a high quality VoWLAN user experience, we strive to minimize jitter in our VoWLAN designs by ensuring that devices have the opportunity to roam well before the quality of the user's voice call is in jeopardy. We accomplish this by assuring that the recommended degree of cell-to-cell overlap exists in our designs.


Figure 5-17 illustrates the concept of cell overlap for a Cisco 7921G VoWLAN handset using 802.11bg. For the Cisco 7921G, the recommended best practices found in the *Voice Over Wireless LAN 4.1 Design Guide* suggest that the cell-to-cell overlap should be approximately 20 percent when using 802.11bg and approximately 15 percent when using 802.11a.

Figure 5-17 20% Inter-Cell overlap



Data applications, on the other hand, typically do not display the same level of sensitivity to packet loss as do voice applications, hence they seldom require the same degree of cell-to-cell overlap. In most cases, a minimum 10% cell-to-cell overlap is sufficient for reliable roaming with data applications, as illustrated in Figure 5-18. High speed data applications and applications combining voice and data capabilities in a single device (smartphones, for example) may require cell-to-cell overlap that resembles a VoWLAN design much more than a data design.

Figure 5-18 10% Inter-Cell Overlap



Although they are possible and do exist, network designs for the location-aware Cisco UWN performed with only location tracking as a use case represent a minority of all Cisco mobility customer installations. Therefore, the designer striving towards completing an optimized location design is likely to be attempting to satisfy the four primary concerns of VoWLAN and data WLAN designers concurrently.

The chief location-tracking concerns of most designers wishing to track asset tags, clients or rogues will center around:

- Perimeter and corner access point placement—Perimeter and corner access point placement is very important to good location accuracy. Refer to [Figure 5-4](#) and the previous discussion surrounding the concept of a “convex hull”. As described earlier, location accuracy tends to fall off the further one strays outside the convex hull encompassing the set of potential device locations on the floor.
- Staggered pattern—Access points should be located on the floor in a staggered fashion to both facilitate an acceptable inter-access point spacing as well bolster the system's ability to perform RSSI multi-lateration for tracked devices.
- Antenna mounting height—In most indoor location applications, antenna mounting height above the area where devices are to be tracked should be ideally between 10 and 15 feet, with 20 feet being a recommended maximum.
- Inter-access point spacing—Access points should be situated so as to minimize any potential risk of degraded location accuracy rises due to:
 - Non-monotonic RSSI versus distance behavior at close range
 - Degradation in the ability of the system to resolve distance based on changes in RSSI.

Generally, this results in access points being deployed with an inter-access point distance of between 40 and 70 feet. However, the coverage requirements of demanding applications (such as voice and high speed data) may require more dense deployments under certain circumstances.

The question that comes to mind then is, can the requirements described earlier for voice and data applications be met in combination with the requirements of location tracking? The answer is yes, with the precise mechanics of how it is done dependent upon the specific requirements of the voice and data applications themselves, the access point and antenna configuration being considered, and the physical characteristics of the environment into which the infrastructure will be deployed. In order to explore this further, we examine the details behind how an access point layout, primarily intended for high speed data and 7921G voice applications, can be further optimized to include location tracking.

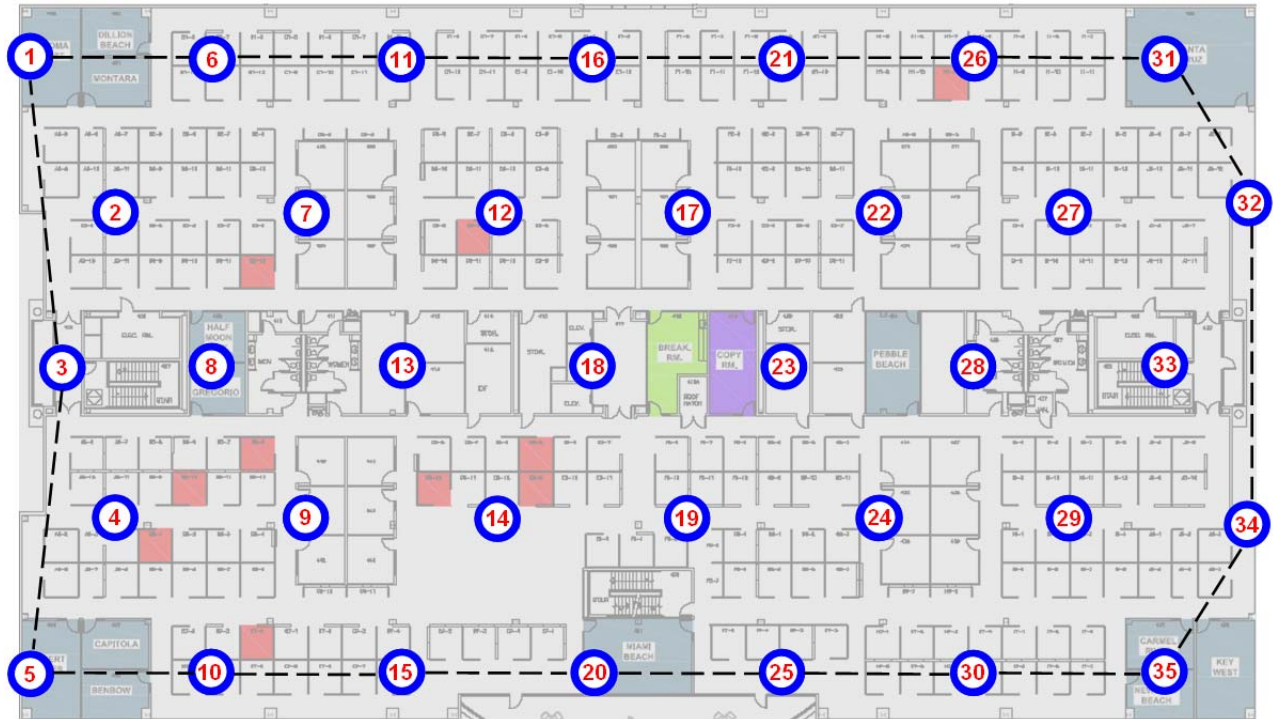
As an example, let us examine a voice access point layout for the 275 x 159 foot facility first presented in [Figure 5-11](#). This represents a drywall office and indoor commercial office environment with a path loss exponent of 3.5 (see [Figure 5-19](#)). These access point locations were selected based on desired signal strength and overlap calculations that were performed by the original designer. In architecting this design, the designer's intention was to provide a solution that closely followed Cisco VoWLAN design best practices described in *Voice Over Wireless LAN 4.1 Design Guide*, which is available at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>

We opted for a dual band infrastructure, with an 802.11a 5 GHz WLAN that is used by 7921G VoWLAN handsets and high-speed WLAN client devices. 802.11bg 2.4 GHz operation is also supported, but due to the substantially reduced overall capacity on 802.11bg brought about by the existence of only three non-interfering channels, its use is restricted to legacy data and voice devices, as well as active RFID asset tags that are in compliance with the Cisco Compatible Extensions for Wi-Fi Tags specification. Legacy data devices would include devices that are unable to migrate to 802.11a for reasons such as the client hardware device being no longer offered for sale, battery life concerns, and so on. Candidate legacy devices might include PDAs, bar code scanners, and other devices with embedded wireless onboard that is not easily upgradeable. In the case of our example, we assume that there are still some

users of 802.11b voice devices present in the environment (for example, legacy Cisco 7920 VoWLAN IP phones, or similar legacy 802.11b-only devices from third parties) that have not yet been addressed with 802.11a replacements.

Figure 5-19 Layout for 5GHz Voice and High Speed Data, 2.4GHz Legacy



In [Figure 5-19](#), we assume the use of 35 ceiling mounted AP1240AG access points, each of which is equipped with a pair of 2.2dBi AIR-ANT4941 antennas for 802.11bg and a pair of 3.5 dBi AIR-ANT5951 antennas for 802.11a. The access points and the antennas are mounted at a height of 10 feet. The design is intended to provide a minimum of -67 dBm signal level and a data rate of at least 24 Mbps on 802.11a for VoWLAN and high speed data clients, and a minimum of -67 dBm signal level and data rate of at least 11 Mbps on 802.11bg for legacy data and voice clients. 802.11a VoWLAN devices are assumed to be Cisco 7921G VoWLAN IP phones with integrated antenna. Legacy voice and data client devices are assumed to possess nominal antenna gain of 0 dBi. Inter-access point spacing is approximately 42.7 feet and was selected to allow for a uniform distribution of access points within the floor interior, and also ensure that the access point power levels required to produce our desired cell-to-cell overlap would fall within the capabilities of our client devices.

Note the following:

- With the exception of access points 1, 5, 32 and 34, access points are not located directly at the floor perimeter. This is not optimal for the support of good location accuracy in all areas of the floor.
- The lack of perimeter access points in the right hand corners of [Figure 5-19](#). Because of this, there are areas in the vicinity of access points 31, 32, 34 and 35 where the location requirement for each point to lie within 70 feet of three different access points in at least three different quadrants (with an access point present in the fourth quadrant at any range) will not be satisfied.
- Transmit power for each access point has been configured to +5dBm for 802.11bg and +11 dBm for 802.11a. This results in a -67 dBm cell radius of approximately 28.72 feet with a cell-to-cell overlap of 15% for 802.11a VoWLAN and high speed data clients. For 802.11bg legacy clients, it results in a -67 dBm cell radius of approximately 31 feet with a 20% cell-to-cell overlap.

**Note**

The transmit power configured for access points should be within the range of the transmit power levels supported by clients to help avoid potential “one-way audio” telephony calls. When using Cisco’s Radio Resource Manager to manage access point power levels, it is further recommended that designers target achieving the required coverage radii and overlap at transmit-power levels that are less than the maximum supported transmit power level of the client device. This is recommended in order to allow the Radio Resource Manager some degree of power allocation “headroom” that can be used to address potential coverage hole situations while still using transmit power levels that are achievable by the client devices.

In order to facilitate optimal location tracking with this design, a few changes, additions and adjustments will be necessary. Examining the current voice and data design and its associated parameters, the current access point spacing, antenna installation height, and placement pattern appear to be acceptable for location usage. However, the lack of access points located at the actual floor perimeters and in the corners of the floor are a concern that should be addressed. This can be seen from the dashed line in [Figure 5-19](#) which illustrates the convex hull established by the current perimeter of access points. Note that areas at each corner and along each upper and lower perimeter lie outside of this boundary. Although these areas may not prove to be a hindrance to some users, for the purposes of this example, our goal is to assure optimal location accuracy in all areas of the floor. This includes the conference rooms in the corners of the floor and in all perimeter areas. Therefore, establishing a proper floor perimeter will be our first order of business.

The first step is to implement top and bottom access point perimeters as close to the building perimeter as feasible, while attempting to maintain the uniform density of access points shown in [Figure 5-19](#) to the highest degree possible. Maintaining a high degree of access point uniformity is especially beneficial to those users that depend on the Cisco Radio Resource Management (RRM) to maintain transmit power control and perform coverage hole remediation. RRM functions most effectively when the distribution of access points on a floor is as uniform as possible.

At this point, we must decide on one of the following options:

1. Expand the equilateral formations comprising our existing access point constellation to accommodate rearranging the top and bottom rows of access points to form the upper and lower portions of the floor perimeter. With this option and our example environment, a minimal number of additional access points would be required, as their primary use is to fill-in any missing areas on the left and right side perimeters. Since it requires expanding the separation between access points, this option is considered more aggressive when compared to option 2 below. Caution must be exercised to avoid modifying the design beyond the limits imposed on access point transmit power (see below).
2. Contract the equilateral formations comprising our existing access point constellation to accommodate shifting upward the current top row of access points and subsequently introducing a sixth row of access points at the bottom to form a new lower perimeter. This option requires a greater number of additional access points when compared to option 1 above. However, since we are reducing the inter-access point distances, this option typically does not possess the risk of increasing access point transmit power levels beyond that of the original design, and is considered the more conservative option of the two.

When considering the first option, it is necessary to examine the current inter-access point spacing and transmit power levels, and estimate the increase that will be required to the inter-access point separation in order to place the existing outer rows of access points at the actual floor perimeters. If current access point transmit power levels are already at high levels relative to the power capabilities of our client devices, and the estimated increase to inter-access point separation appears to be large, then expanding the constellation of existing access points to accommodate perimeter placement may not be the best


option. This is mainly because it may require the use of higher than desirable access point transmit power levels. In such cases, it is recommended to pursue the second option, which contracts the equilateral formations and results in shorter inter-access point separation, typically with the same or reduced access point transmit power levels.

Recall from our discussion that our transmit power levels are configured at +5dBm for 802.11bg and +11 dBm for 802.11a. In order to determine the new inter-access point separation that would be in effect if we were to uniformly expand the current formations (seen as equilateral triangles in Figure 5-19), we need to perform some basic geometrical calculations. We determine the new inter-access point separation required by assuming that the current top and bottom rows of access points are relocated such that they are positioned at the actual top and bottom floor perimeter. For the 275 x 159 floor in Figure 5-19, this is performed by dividing the top-to-bottom width of the floor (159 feet) by the number of desired rows of equilateral triangular formations (4), thereby yielding a projected formation height of 39.75 feet.

From the premise that in an equilateral triangle each angle is equal to 60° (shown in Figure 5-20), we calculate the length of any side s from the height h of our equilateral triangle formations as follows:

$$h = s(\sin 60^\circ)$$

Figure 5-20 Equilateral Access Point Formation



Solving for s , we calculate $s = \frac{h}{\sin 60^\circ}$ or $\frac{39.75}{.866} = 45.9$ feet. Thus, we would need to expand our current inter-access point spacing from 42.7 feet to 45.9 feet in order to move both the top and bottom rows of outermost access points to the actual building perimeter. As this represents a relatively minor increase in inter-access point spacing, it should be easily accommodated by a correspondingly minor increase in transmit power, if any at all. In our next step, we determine the new cell size that would be required to support the recommended levels of overlap, given our newly calculated inter-access point spacing.

Using this new value for inter-access point spacing, we first calculate the -67dBm cell signal boundary with a 15% cell-to-cell overlap for 802.11a. We then calculate the -67dBm cell signal boundary with a 20% cell-to-cell overlap for our legacy data and voice devices that will be using 802.11bg. With the assumption that the radii of any two adjacent access point cells are equal (that is $R_1=R_2=R$), we can use the equation for the area of a circle-circle intersection as the basis for this calculation. To determine the cell radius given that the inter-access point separation and the percentage of overlap are known, we proceed as follows:

$$O\pi R^2 = 2R^2 \arccos\left(\frac{d}{2R}\right) - \frac{1}{2}d\sqrt{4R^2 - d^2}$$

where:

- O = the desired overlap percentage divided by 100
- $\arccos\left(\frac{d}{2R}\right)$
- $\arccos\left(\frac{d}{2R}\right)$ is expressed in radians
- d = the inter-access point distance in feet
- R = the cell radius in feet

We substitute either 15 (for 802.11a) or 10 (for 802.11bg) as the percentage of overlap O , and 45.9 feet for the inter-access point distance d . Solving for R as an approximate root of the function shown above, we determine that the cell radii should be equal to 30.88 feet for a 15% cell-to-cell overlap using 802.11a and 33.4 feet for a 20% cell-to-cell overlap using 802.11bg¹.

At this point, we have the information necessary to calculate the access point transmission power settings that will be necessary to achieve our desired cell signal boundaries. This can be performed using a form of the equation presented earlier to calculate receive signal strength (RX_{POWER}) from knowledge of our reference path loss, path loss exponent, transmit power and various miscellaneous receive and transmit gains and losses. This was discussed in [Received Signal Strength \(RSS\), page 2-7](#). As it is the transmit power (TX_{POWER}) of our access points that we wish to calculate and not the receive signal strength, we shall use a modified form of the equation as follows:

$$TX_{POWER} = RX_{POWER} + LOSS_{TX} - Gain_{TX} + PL_{METER} + 10\log D^n + s - Gain_{RX} + LOSS_{RX}$$

For the purposes of this example, we have assumed:

- That transmission losses due to cables, connectors, etc ($LOSS_{TX}$ and $LOSS_{RX}$) are equal to 0 dB.
- 0 dB shadow fading standard deviation.
- Receive antenna gain for our legacy 2.4 GHz data client devices of 0 dBi.

Substituting the appropriate values along with our expectation of a -67 dBm minimum receive signal strength (RX_{POWER}) for both 802.11a 802.11bg, as well as the appropriate antenna gains, our cell radius in meters (30.88 feet = 9.41 meters, 33.4 feet = 10.18 meters), an estimated path loss exponent n of 3.5 and our reference path losses, we obtain the following results:

802.11bg:

$$\begin{aligned} TX_{POWER} &= -67 \text{ dBm} + 0 - 2.2 \text{ dBi} + 40 \text{ dB} + 10\log(10.18^{3.5}) - 0 + 0 \\ &= -29.2 + (10 * 3.527) \\ TX_{POWER} &= +6.07 \text{ dBm, or approximately } +8 \text{ dBm} \end{aligned}$$

802.11a:

$$\begin{aligned} TX_{POWER} &= -67 \text{ dBm} + 0 - 3.5 \text{ dBi} + 46 \text{ dB} + 10\log(9.41^{3.5}) - (-3.0) + 0 \\ &= -24.5 + (10 * 3.408) + 3 \\ TX_{POWER} &= +12.58 \text{ dBm, or approximately } +14 \text{ dBm} \end{aligned}$$

1. A computer algebra system (CAS) capable of both symbolic and numeric calculations, such as Maple, Mathematica or Maxima was found to be helpful in solving such calculations. See [Appendix B](#) of this document for information regarding how to use Maxima to calculate R as an approximate root of the aforementioned equation over the closed interval ($d/2$, d).

Note that these power levels have been rounded upward to the next available transmit power increment available on the AP1240AG access point. Since this is +1.93 dBm higher than the required transmit power to achieve our recommended 20% overlap goal at a cell signal boundary of -67 dBm, we can expect that the overlap will exceed the 20% target. This is acceptable, as the 20% overlap is a minimum target. Similarly, for 802.11a the access point transmit power level of +14 dBm is +1.42 dBm higher than what is required to achieve the recommended 15% overlap, once again resulting in more overlap between cells than expected.

In this particular case, the option to expand our inter-access point separation is an acceptable alternative. Due to the increase in the inter-access point separation (from 42.7 feet to 45.9 feet), a +3 dBm increase is required to both our 802.11a and 802.11bg access point transmit power settings in order to remain in strict compliance with our calculated requirements. Despite the increase in access point transmit power level, additional transmit power is left in reserve on both bands to address potential coverage holes or other anomalies that could occur due to changes in the environment. If this had not been the case, we would have proceeded with our second option which entails contracting our inter-access point spacing and introducing a sixth row of access points. The main differences in our calculations would be to divide the size of floor by five (instead of four) rows of equilateral triangular formations. This would have resulted in a smaller formation height, a smaller inter-access point separation, and therefore, smaller cell-to-cell radii and lower transmit powers.

**Note**

The signal level measurements and the calculations described in this section, while based on generally accepted RF theory, are intended for planning purposes only. It is reasonable to expect some level of signal level variation from these theoretical calculations in different environments.

Rather than statically administering access point transmission power levels, the Cisco Radio Resource Manager (RRM) can be used instead. RRM can be used to dynamically control access point transmit power based on real-time WLAN conditions. Under normal circumstances, transmit power is maintained across all access points to maintain capacity and reduce interference. If a failed access point is detected, transmit power can be automatically increased on surrounding access points to fill the gap created by the loss in coverage. Should a coverage hole occur, RRM can use any remaining transmit power reserve on surrounding access points to raise the adjacent coverage levels and address the coverage hole until it can be investigated and resolved.

In either case, it is recommended that a verification of access point transmit power settings be performed periodically. If you opt to manually administer access point transmit power settings, you should examine the overall performance of your system to ensure that your original design assumptions are still valid and that there have not been significant changes in your environment that might warrant reconsideration of those assumptions. When using RRM, it will monitor your system for changes that might warrant an increase or decrease in access point transmit power settings for you. After your system has been installed, various adjustments can be made to RRM to bring its selection of access point transmit power levels and other parameters within your expectations for the environment at hand.

Keep in mind that immediately after installation and for a period of time after, it is reasonable to see a fair degree of RRM activity, as the system settles in and final parameter selections are made. At the conclusion of this “settling in” period, the system designer should ensure that the choices made by RRM are inline with the overall expectations of the design. Once the system has settled there should be little to no change in RRM managed parameters over time, as barring any significant environmental or equipment changes, the selections made for access point transmit power levels should remain fairly static. Any indication of constant fluctuation in assigned access point transmit power levels or channels should be regarded by the system administrator as potential indication of other anomalies that may be developing within the environment. The root causes behind such frequent fluctuations should be investigated and addressed promptly.

**Note**

A comprehensive discussion of the mechanics of RRM is beyond the scope of this document. For information of this nature, it is highly recommended that readers refer to *Radio Resource Management under Unified Wireless Networks* document, which can be found at the following URL: http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml. In addition, it is recommended that all users considering using RRM in VoWLAN designs refer to the *Voice Over Wireless LAN 4.1 Design Guide*, which can be found at the following URL: <http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>

Figure 5-21 illustrates the updated access point layout, using the information from the calculations above along with perimeter access point placement which is discussed next.

Figure 5-21 Layout for 5GHz Voice and High Speed Data, 2.4GHz Legacy, with Location



In Figure 5-21 we can see the effects of the increase in inter-access point distance:

- The top row (access points 1, 6, 11, 16, 21, 26 and 31) and bottom row (access points 5, 10, 15, 20, 25, 30, and 35) of access points are now located at the actual top and bottom floor perimeter.
- On the right side of the floor perimeter, access points 31 and 35 have been moved into the right hand corners of the floor. Access point 33 has been moved to the right side of the floor perimeter. As a group, access points 31 through 35 now comprise the right side of the floor perimeter.
- On the left side of the floor perimeter, access points 1 and 5 have been moved into the left hand corners of the floor. In addition, two new access points (36 and 37, indicated by adjacent yellow stars) have been added to the design to complete the formation of the left side of the floor perimeter.

The two new access points added in Figure 5-21 bring the total access point count for the integrated voice, data and location design to 37 access points. The primary source of voice and data coverage in this design still emanates from the access points participating in the equilateral formations seen across

223339

the floor (i.e. this can be seen in [Figure 5-21](#) as the set of access points depicted in red). Access points 32, 34, 36 and 37 are necessary to establish a location perimeter, but based on the assumptions and calculations presented here, may not be required to participate in providing voice or data coverage in either band. That being the case, these access points can be statically configured to operate at significantly reduced transmit power (such as -1 dBm, for example), which also minimizes the co-channel interference contribution of these access points as well.

**Note**

For information regarding co-channel interference concerns in VoWLAN designs, it is recommended that readers refer to the *Voice Over Wireless LAN 4.1 Design Guide*, which can be found at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/vowlan/41dg/vowlan41dg-book.html>.

When using Cisco RRM to manage power levels, access points that are placed into the design solely for location purposes should *not* be included in either the Radio Resource Management transmit power control or coverage hole remediation processes. Configuring a custom access point transmit power level (using the “custom” TX power option on WCS or the controller GUI) will automatically exclude these access points from transmit power and coverage hole remediation algorithms.

Based on our planning output and our calculations, our original voice and data design shown in [Figure 5-19](#) can be migrated to a location-ready design that is in compliance with the best practices described in the *Voice Over Wireless LAN 4.1 Design Guide* with only minor changes in both layout and configuration. The result is a combined design that is well suited to support VoWLAN, high speed data and location tracking on 5 GHz, as well as legacy data and voice support with location tracking on 2.4 GHz.

Additional activities that can be performed to improve designs and design implementation include:

- Performing a walk-around of the site and verifying that areas on the floor plan where access point mounting is desired can actually accommodate it. This is always a good idea, since floor plans and blueprints do not always indicate the precise conditions present at each location where an access point may be mounted. For example, you may find that certain locations that appear to be viable candidates “on paper” actually are inaccessible (such as an electrical closet), inappropriate (such as an outdoor balcony) or are otherwise not acceptable. In such cases, access points should be relocated close to the original location such that the impact on the overall design is minimal. In [Figure 5-21](#), some common-sense obstacles have been avoided, and the affected access points have been moved slightly.
- Verifying RF propagation and coverage assumptions by temporarily installing a few access points in various test areas of the floor, and measuring actual RF signal strength and cell-to-cell overlap using a portable client device with appropriate “site survey” software tools. This is an excellent time to measure the ambient noise levels of the potential access point cells as well, and determine whether the projected signal to noise ratio will be sufficient. Note that Cisco's RRM feature also monitors client SNR and increases access point power if a number of clients are noticed to fall below a prescribed SNR threshold. For more information about RRM, refer to the *Radio Resource Management under Unified Wireless Networks* at the following URL:
http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml
- Validating whether there are any radar users present in your locale that may interfere with the use of the additional 802.11a that are subject to DFS. If there are not, these channels can be made available for use by enabling DFS on your WLAN controllers.

**Note**

In software Release 4.1 of the Cisco Unified Wireless Network, Cisco Compatible Extensions Location Measurements are not enabled for any 802.11a channels on which DFS operation has been mandated. For DFS channels, RSSI from probe requests transmitted by WLAN clients as a part of their normal operation will be used for location tracking purposes.

The techniques and principles described in this section illustrate how a design performed in accordance with VoWLAN and data best practices can be upgraded to being “location-ready”. The key concepts behind how inter-access point separation, cell radius and transmit power are inter-related, and how these factors can be used to determine coverage overlap, can be applied to designs of various different sizes and shapes, as well as environments with varying path loss characteristics and shadowing.

Avoiding Location Display Jitter

Location smoothing was introduced to enable the network administrator to compensate for cases of location instability sometimes seen with clients that are not actually experiencing any change in movement. This can be due to a variety of factors, including the following:

- Variations in client transmit power resulting in changes in detected RSSI
- Environmental changes, including semi-permanent obstructions that may have shifted position, result in variations in attenuation and multi-path
- Changes in client orientation
- Shadow fading

Location smoothing allows for varying degrees of averaging to be applied to device location. Smoothing factors are set in **Location > Location Server > Administration > Location Parameters** through the *Smooth Location Positions* parameter, as shown in [Figure 5-22](#).

Figure 5-22 Configuring Location Smoothing

Location Server > Location Parameters > 'AeS_Loc2'

Location Parameters

Enable calculation time [?](#) Enable

Enable OW Location [?](#) Enable

Relative discard RSSI time [?](#) minutes.

Absolute discard RSSI time [?](#) minutes.

RSSI Cutoff [?](#) db.

Smooth Location Positions [?](#)

Off (no smoothing)

Less smoothing (new value weighted more)

Average smoothing (new value weighted same)

More smoothing (new value weighted less)

Maximum smoothing (new value weighted min)

190572

The various smoothing factor options impact the displayed location position by assigning different weights to the latest calculated position of the device versus its last known position. These weights are assigned as shown in [Table 5-1](#).

Table 5-1 Smoothing Factor Weight Assignments

Smooth Location Positions Value	Weight Assigned to Previous Position	Weight Assigned to New Position
Off (no smoothing)	0%	100%
Less smoothing	25%	75%
Average smoothing	50%	50%
More smoothing (default)	75%	25%
Maximum smoothing	90%	10%

As the weight assigned to the previous position is increased in relation to the weight assigned to the new position, the amount of displayed device movement is decreased. Note that the use of location smoothing will not eliminate all observed movement in the location display for that device. Rather, the use of location smoothing simply limits the rate at which such changes are communicated to the end user.

The use of location smoothing involves a small tradeoff between location viewing stability and the reaction time of the location display to changes in position. For most environments, the use of the default smoothing factor should provide an improved viewing experience. Higher smoothing factors are best reserved for environments where there is very infrequent movement of WLAN clients and tagged assets. Low smoothing factors (or no smoothing) may provide better results in situations where tagged assets and clients are frequently moving, and in some cases experiencing constant or near-constant motion. Users that are primarily concerned with minimizing displayed location latency (albeit at the risk of some location jitter) should choose low values for location smoothing, or disable location smoothing altogether.

Multiple Location Appliance Designs

As stated earlier, a single Cisco Wireless Location Appliance can track up to 2500 devices, which includes WLAN clients, asset tags, rogue access points, and rogue clients. The location appliance allows for specific tracked device categories to be enabled via **Location > Location Server > Administration > Polling Parameters**. To make best use of the capacity of each location appliance, Cisco recommends enabling only those polling categories (client stations, rogues, asset tags, or statistics) in which there is genuine interest and that require simultaneous tracking/historical location. For example, if the primary interest is in tracking asset tags only, do not enable the client and rogue polling categories because this only adds to overall network traffic between the location appliance and WLAN controllers as well as unnecessarily consuming a portion of the 2500 device tracking capacity. By disabling polling for device categories for which there is little interest, the full capacity of the location appliance can be better used.



Note

Although not the focus of this document, Release 4.2 of the location-aware Cisco UWN introduces an enhancement that allows for individual limits to be placed on what portion of the location appliance's entire tracked device capacity is allocated to each tracked device category (i.e., WLAN clients, asset tags, and rogue devices).

- Cisco WCS Release 4.1 can support between 500 and 3000 access points and between 125 and 750 WLAN controllers, depending on hardware configuration, as follows:
- WCS high-end server—Supports up to 3000 access points and 750 WLAN controllers
- WCS standard server—Supports up to 2000 access points and 500 WLAN controllers
- WCS on Cisco Wireless LAN Solutions Engine (WLSE) hardware—Supports up to 1500 access points and 100 WLAN controllers
- WCS low-end server—Supports up to 500 access points and 125 WLAN controllers

**Note**

For complete details, see the *Cisco Wireless Control System Configuration Guide, System Requirements* at the following URL:

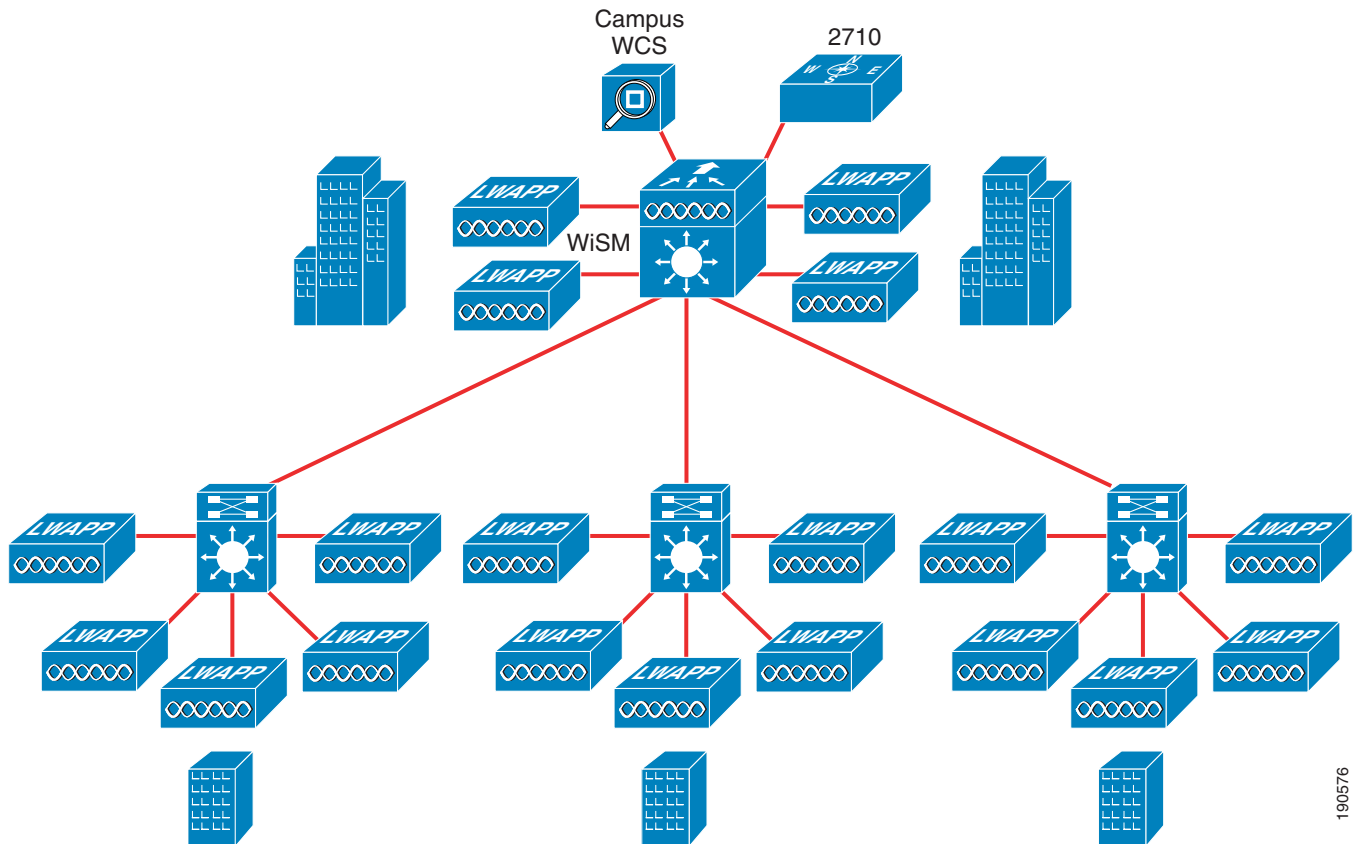
http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_chapter09186a00808317da.html#wp1061082

The maximum size of the WCS *management domain* (that is, the total number of devices managed by a single WCS) is related to the choice of server platform. In very large networks, it may be necessary to partition the network into multiple management domains, each with a separate WCS managing it. Beginning with software Release 4.1, Cisco introduces the WCS Navigator, which is a management aggregation platform for the enhanced scalability, manageability, and visibility of large-scale implementations of the Cisco Unified Wireless Network. WCS Navigator is a software-based solution that enables overall management of multiple Cisco WCS management platforms, regardless of their physical location. Cisco WCS Navigator runs on a separate server platform with an embedded database, and supports up to 20 Cisco WCS management platforms and up to 20,000 access points. Additional information regarding Cisco WCS Navigator can be found at the following URL:

<http://www.cisco.com/en/US/products/ps7305/index.html>.

Although the maximum size of the each management domain is related to the capacity of the WCS platform managing it, the maximum size of the *location domain* (that is, the number of devices tracked by a single location appliance) is limited by the tracked device capacity of the location appliance. In a large percentage of mid to large size deployments, the standard deployment model of a single WCS management domain combined with a single location domain (Figure 5-23) can meet the device tracking and management needs.

Figure 5-23 Single Management and Location Domains



190576

However, in the case of larger campuses, the total number of tracked devices may exceed the capacity of a single location appliance, making it necessary to use multiple location appliances (and multiple location domains) in the design. Some organizations may in fact choose to purposely divide the tracked device load among two or more location appliances for internal reasons, such as to better accommodate internal cost accounting within the organization, or to better accommodate growth. A good example of this might be a campus medical center WLAN that is tracking a large amount of patient-related medical assets in addition to the internal IT assets of the organization. It may wish to use separate location appliances to partition the tracking of assets as well as to provide clear delineation of equipment and cost separation between departments.

**Note**

Keep in mind that if there are more than 250 tracked devices detected in a particular device category, WCS displays only the first 250 tracked devices. After displaying the initial 250 tracked devices, WCS will prompt the user to initiate the use of device filtering in order to display the remainder.

In software Release 4.1, WLAN controllers offer the device capacities shown in [Table 5-2](#).

Table 5-2 *WLAN Controller Device Capacities*

Controller Model	WLAN Clients Supported	Asset Tags Supported	Rogue APs Supported	Rogue Clients
2006	256	500	125	100
2106	256	500	125	100
4402	2,500	1250	625	500
4404	5,000	2500	625	500
WiSM	10,000	5000	1250	1000
NM-WLC6	256	500	125	100
NME-WLC8/12	350	500	125	100
3750G	2,500	1250	625	500

The subsections that follow examine how WCS, the location appliance and WCS Navigator can be combined to satisfy the needs of more demanding designs that are typically beyond the capabilities of single management/location domain combinations.

Single Management Domain with Multiple Location Domains

In this design approach, the network management needs of the enterprise WLAN are expected to be well within the capacity of a single WCS management domain. However, there is a need to track a combination of more than 2500 clients, rogues, and asset tags. This can be accomplished using multiple location domains that are managed via a single management domain, as shown in Figure 5-24.

Figure 5-24 Single Management Domain with Multiple Location Domains

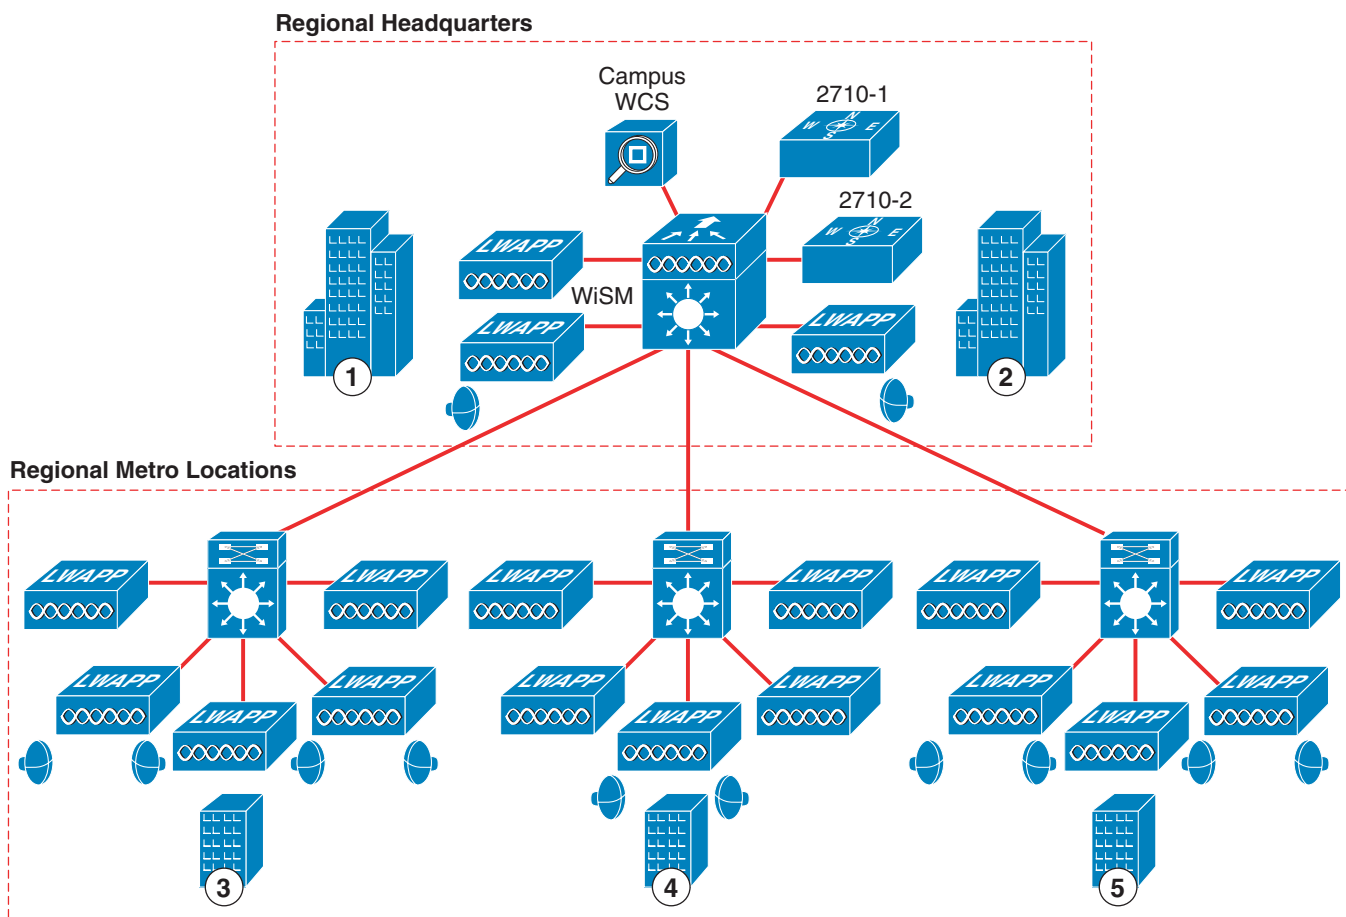


Figure 5-24 illustrates a single campus WCS server providing WLAN management services for a large regional headquarters campus location as well as three extended regional metropolitan campus locations. In this case, all are located within a major metropolitan city.



Note

Although the locations in this example are geographically dispersed, the concepts discussed regarding multiple location domains and controller assignments also applies to the case of a single contiguous campus where the total number of tracked devices exceeds 2500.

In the example depicted in Figure 5-24, the regional headquarters location contains 140 access points and each metro location contains 50 access points along with an unspecified number of chokepoint triggers. The design calls for the use of a centralized Cisco Catalyst 6500 with Wireless Service Module (WiSM). The WiSM contains two embedded controllers per service module, which are referred to as WiSM-1 and WiSM-2. WiSM-1 is used to service access points at the regional headquarters location while WiSM-2 services access points located at the metropolitan locations. Two Cisco 2710 Wireless

Location Appliances provide the capacity to track up to a maximum of 5000 device MAC addresses. Location appliance 2710-1 is assigned to track assets within the regional headquarters complex only, and location appliance 2710-2 tracks assets across all three of the metropolitan locations. Note that location appliance 2710-2 does not track devices in the regional headquarters complex.

We can use WCS to create a single network design (a set of outdoor, campus, building and floor maps along with access point and chokepoint placements) that encompassing the entire extended campus shown in [Figure 5-24](#). We first add in buildings 1 and 2 and the floors that are included in each building for the regional headquarters location. The 140 access points that are registered to controller WiSM-1, along with any chokepoints that are in use at the regional headquarters, are assigned to this network design. In addition, an event notification group is created for the regional headquarters location. WCS is then used to add metro remote buildings 3, 4, and 5 and their respective floors. The 150 access points that are registered to controller WiSM-2, along with any chokepoint triggers in the metro remote locations, are assigned to this network design. A separate event notification group is created for the metro remote locations.

The critical step in this process is not only to share the network design between both location appliances, but to ensure that the WLAN controllers we wish included within the location domain of each location appliance are correctly synchronized to (and *only* to) that location appliance. Thus, while our campus network design is synchronized to both location appliance 2710-1 and 2710-2, *only* controller WiSM-1 is synchronized to location appliance 2710-1 and *only* controller WiSM-2 is synchronized to location appliance 2710-2.

After these actions are performed, we will be able to manage the entire campus from the single WCS management domain, while dividing the aggregate number of tracked devices between two separate location domains. Once implemented, the entire enterprise is managed as a single management domain, with all management polling and reporting emanating from a centralized WCS. Location appliance 2710-1 handles polling controller WiSM-1 for all information pertaining to tracked devices found within its location domain, which is the regional headquarters. Location appliance 2710-2 handles the polling of controller WiSM-2 with regard to all tracked devices found in its location domain, which are the regional metro locations. Except for the fact that the two location domains operate across a common network, are managed from a common management domain, and possesses a controller that co-resides on the same physical WiSM module at the regional headquarters, the two location domains essentially exist independent of one another.

An alternate approach would entail creating two separate location appliance network designs containing only the portion of the campus network encompassed within each respective location domain. Each of these individual network designs would then be synchronized to the appropriate location appliance, along with the controller(s) that the location appliance will service. While this approach is indeed functional, in general, the single network design approach is preferred from a convenience and ease of maintenance perspective. On the other hand, when working with very large network designs, the dual network design approach may be preferable. By splitting the environment into two distinct network designs, it avoids loading the details about unnecessary access points, buildings, floors and the graphical images representing them into location appliances that have no need for this information.

In any design containing two or more independent location domains, it is important to be aware of the degree and frequency of inter-domain tracked device migration, if any. Some designs involving two or more location domains are meant to operate as pure “closed loop” systems. This implies that tracked devices are added and removed under very controlled circumstances, with the exception of any newly discovered rogues. In a closed loop system, the number of non-rogue tracked devices in each location domain can be expected to remain fairly constant. However, many actual deployments may employ “open-loop” business processes, where some degree of uncontrolled device addition, removal or migration can, and often will occur.

This can be significant since tracked devices that migrate from one location domain to another may appear in the active location database of each location appliance until such time that they are pruned. A tracked device will be pruned from the active location database of a location appliance when all of the following conditions are true¹:

1. The device is no longer being detected by any of the access points registered to any controller being serviced by that location appliance (i.e., any controller that has been synchronized to that location appliance).
2. All controllers whose access points have detected the tracked device have ceased reporting RSSI for the device to the location appliance (typically after the controller's RFID timeout or client user timeout has expired).
3. the location appliance Absent Data Cleanup Interval (ADCI) has expired (default is 1440 minutes).

During the time that a tracked device is present in more than one location appliance, multiple entries for the same device may appear in WCS client, tag or rogue display menus. An example of this is depicted in Figure 5-25, which illustrates the results of an asset tag search across all location appliances in lab testing. In this test, an asset tag with MAC address 00:0C:CC:5E:82:90 has migrated from the domain of location appliance "Loc1" to the domain of location appliance "Loc2". The access points of the two buildings are registered to different controllers, and these controllers are individually contained within two different location domains. For example, the controller servicing the hypothetical Havermeyer Building is being polled by the "Loc1" location appliance while the controller servicing the hypothetical Pupin Hall Building is being polled by the "Loc2" location appliance.

Figure 5-25 Duplicate Device Appearances Due to Device Migration

Tags

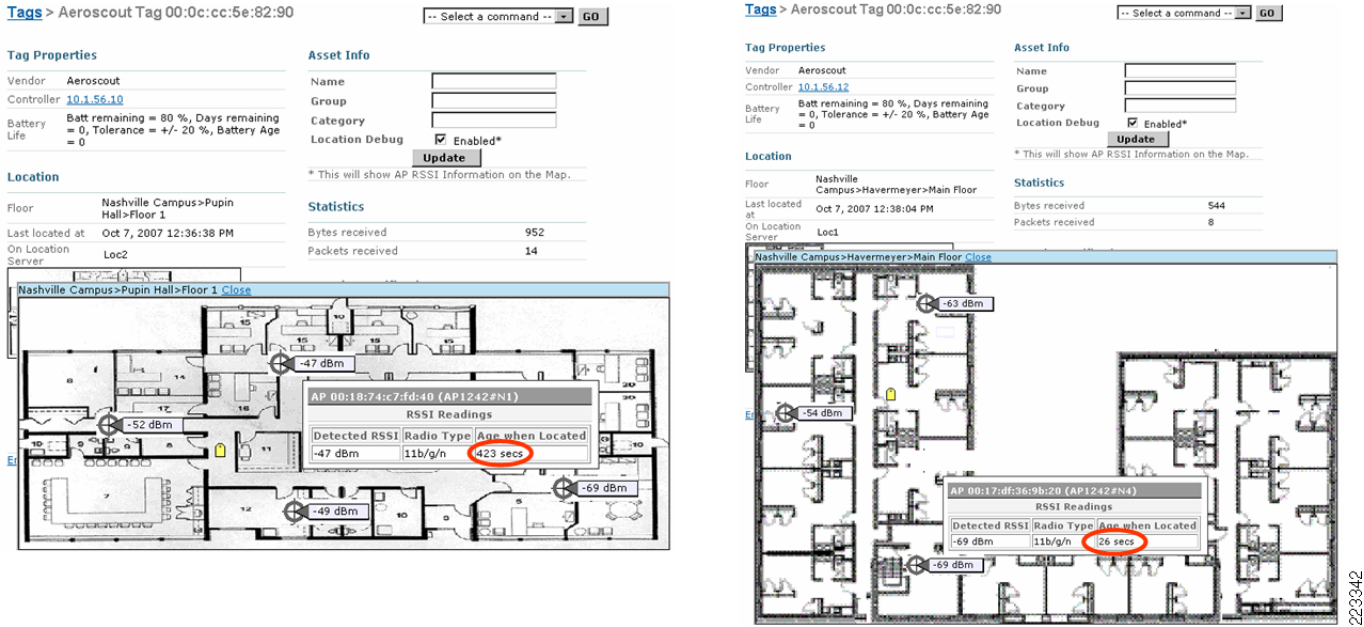
MAC Addr	Asset Name	Asset Group	Asset Category	Vendor	Loc Server	Controller	Battery Status	Map Location
00:0c:cc:5e:82:90	-	-	-	Aeroscout	Loc1	10.1.56.12	80 %	Nashville Campus>Havermeyer>Main Floor
00:0c:cc:5e:8d:cc	-	-	-	Aeroscout	Loc2	10.1.56.10	80 %	Nashville Campus>Pupin Hall>Floor 1
00:0c:cc:5e:82:90	-	-	-	Aeroscout	Loc2	10.1.56.10	80 %	Nashville Campus>Pupin Hall>Floor 1

223941

Note that in Figure 5-25, the asset tag MAC address 00:0C:CC:5E:82:90 appears twice, found in two different location appliances and with locations listed for two different buildings. If we were to display the map locations listed at the extreme right of Figure 5-25, we would see a tag icon displayed on both the main floor of the Havermeyer Building as well as floor number one of Pupin Hall. In order to differentiate between which of these two entries represents the latest and true location of asset tag 00:0C:CC:5E:82:90, we can use the information provided by the "Last Located At" time stamp and the age of the RSSI readings used to locate the device. To obtain the RSSI age information, we click on each appearance of device MAC address 00:0C:CC:5E:82:90, enable the location debug option on the tag detail screen that appears, and view the age of the RSSI readings used to establish the device location (shown in Figure 5-26).

1. The length of the location appliance's polling cycle for the particular category of device may also play a role. The precise degree of impact will be dependent on the point within the polling cycle the tracked device was prior to migrating out of the location domain.

Figure 5-26 RSSI Age Comparison for Duplicate Tag Entries



After the location appliance’s absent data cleanup interval has expired, our lab test tag report appears as shown in Figure 5-27. Note that our migrated asset tag with MAC address 00:0C:CC:5E:82:90 has been pruned from the active location database of location appliance “Loc1”.

Figure 5-27 Tag Report After Device Pruning

Tags

MAC Addr	Asset Name	Asset Group	Asset Category	Vendor	Loc Server	Controller	Battery Status	Map Location
00:0c:cc:5e:8d:cc	-	-	-	Aeroscout	Loc2	10.1.56.10	80 %	Nashville Campus>Pupin Hall>Floor 1
00:0c:cc:5e:82:90	-	-	-	Aeroscout	Loc2	10.1.56.10	80 %	Nashville Campus>Pupin Hall>Floor 1

Relating this to the example shown in Figure 5-24, devices moving between metro remote buildings 3, 4, and 5 would not be affected by this anomaly, since these buildings are all contained within the location domain of appliance 2710-2. However, it would be of concern with devices migrating between any of the metro remote buildings and regional headquarters buildings 1 and 2, since the regional headquarters buildings are contained within a different location domain.

In cases where tracked devices may migrate between location domains at a more or less equal rate (that is, the rate of devices leaving a location domain is approximately the same as that of those entering) and the absent data cleanup interval is left at the system default, the situation illustrated above may persist for approximately 1440 minutes (24 hours). If a significant degree of device migration is expected, it is a good idea to tune the absent data cleanup interval to the anticipated level of migration expected at your site. This will help make more efficient use of location appliance resources and help avoid the possibility of exhausting tracked device capacity.

For example, assume a situation where a particular site begins its workday with:

- Its location appliance's tracked device capacity at 60% (1500 devices).
- A bi-directional two device per-minute migration rate (clients and tags).
- No other tracked device additions or removals (rogue tracking is disabled).

In this case, the remaining capacity on the location appliance (1000 tracked devices, also referred to as “headroom”) would allow for 500 minutes of operation at this rate of device migration before the potential for tracked device capacity exhaustion becomes a concern. To help avoid this situation, the Absent Data Cleanup Interval should be set *below* 500 minutes (for example, 440 or 470 minutes) in order to reclaim tracked device capacity prior to headroom depletion.

A common question that often arises is “why not just set the Absent Data Cleanup Interval to a very short interval initially?” The answer revolves has to do with decreases in the Absent Data Cleanup Interval not being entirely without tradeoffs:

- While reducing the location appliance's Absent Data Cleanup Interval to an arbitrarily short value may have the effect of freeing up device capacity quickly, it can also effect the ability to view prior location and statistics for all categories of tracked devices (asset tags, clients and rogues) for which updated RSSI information has not been received within the ADCI time period. If the ADCI has been reduced and a device has been removed from the active location database, the ability to view past location information via the “Load Location Data as Old As” dropdown menu in the left hand margin of WCS floor maps display screens will now be limited to the reduced length of the Absent Data Cleanup Interval. Normally, this function allows information to be extracted from the active location database for up to a full 24 hour period, matching the default time period of the Absent Data Cleanup Interval. If the Absent Data Cleanup Interval is reduced, the maximum scope of the “Load Location Data As Old As” dropdown is reduced to the new value for the Absent Data Cleanup Interval, despite values higher than this being displayed in the dropdown selector.
- When devices are pruned from the active location databases, historical location information on that location appliance will not be accessible for those device MAC addresses until such time that these devices re-enter the location domain and are re-added to the active location database. This assumes that the historical location information has not been pruned via an the independent location history pruning process that occurs periodically.

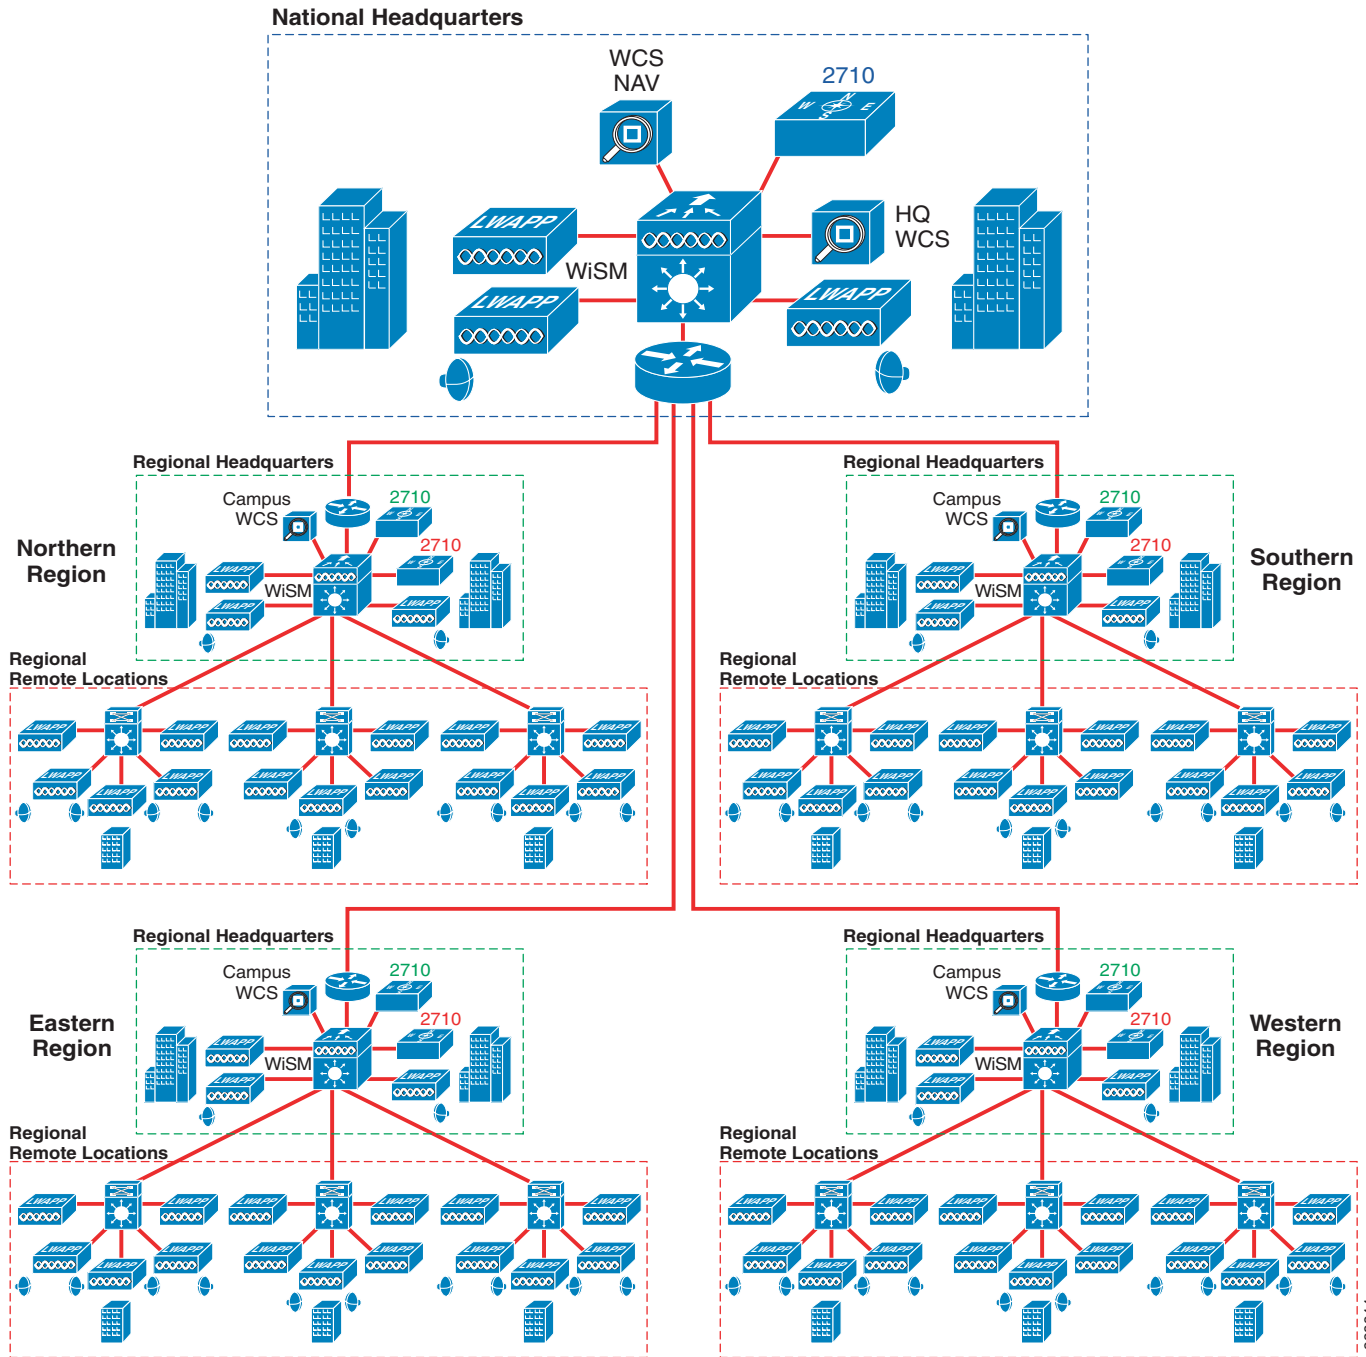
Multiple Management Domains with Multiple Location Domains

In this section, we examine the case where both of the following are true:

- The combined number of access points and controllers managed within the enterprise cannot be contained within a single management domain
- The number of tracked devices in the enterprise exceeds the capacity of a single location domain.

We revisit the enterprise depicted in [Figure 5-24](#), except we now examine the organization's structure from a perspective higher up in the organization's hierarchy. We see that the headquarters location and the metropolitan remote office extended campus locations that we discussed in [Figure 5-24](#) actually comprise a regional entity, with the headquarters location in [Figure 5-24](#) now representing a “regional” headquarters location. Each regional headquarters reports to a national headquarters location, and there are three other regional headquarters locations that are very similar to what we described in [Figure 5-28](#). In other words, we see that the “enterprise” we discussed in [Figure 5-24](#) is really part of a much larger entity. In this section, this larger entity is discussed in detail using [Figure 5-28](#), with much of our discussion building on the information we covered in the previous section and [Figure 5-24](#).

Figure 5-28 Multiple Management Domains with Multiple Location Domains



Each of the four regions depicted in Figure 5-28 contains a total of 290 access points, with a WCS server resident at each regional headquarters. As we discussed in the previous section, two location appliances are deployed per region and physically reside within each regional headquarters location, managing the two location domains present within each region. A national headquarters campus provides overall management for the entire enterprise, and contains 200 access points and a total of 2,100 tracked devices

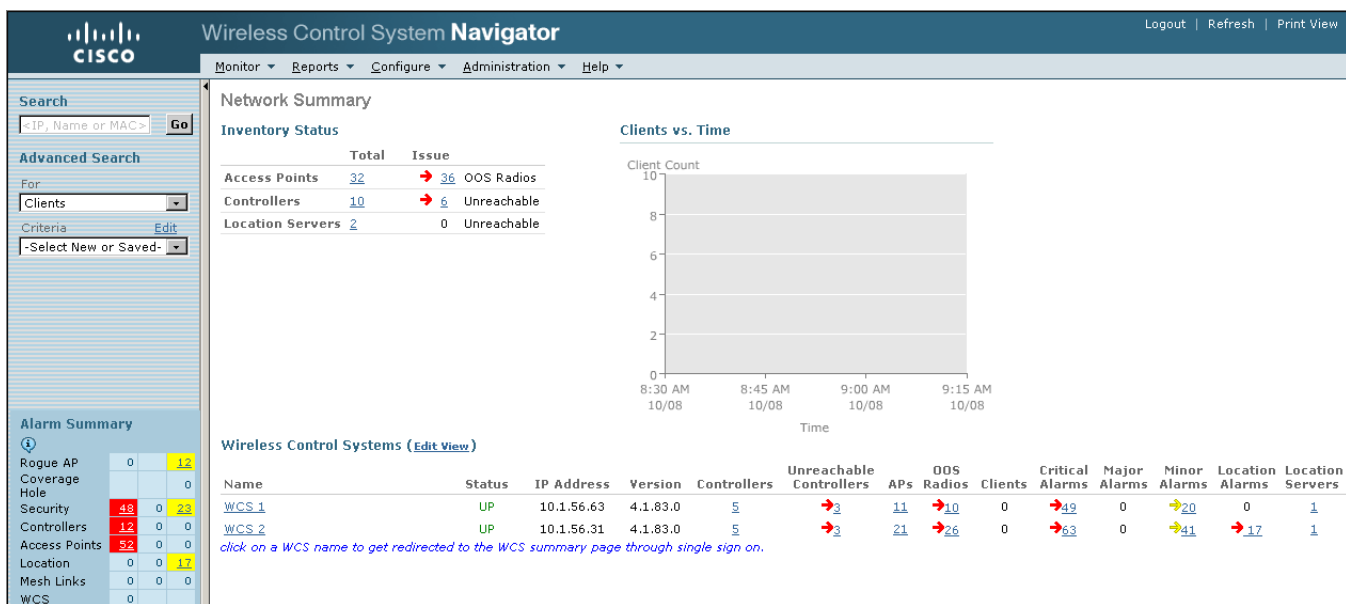
223344

of its own, with WLAN controller services provided by a resident WiSM-equipped Catalyst 6500. The national headquarters location contains a single management and a single location domain, primarily dedicated to the WLAN and location tracking needs of the national headquarters staff itself.

This design contains a total of nine location domains and five management domains. To provide top-down management of the entire Unified Wireless Network, WCS Navigator is deployed at the national headquarters location, providing both management and location visibility internally as well as to all four regional campuses. Each WCS server communicates with WCS Navigator via its northbound API, allowing for the events, activities and resources of each regional management domain to be monitored and managed via a centralized portal. Figure 5-28 contains a single management and a single location domain, primarily dedicated to the WLAN and location tracking needs of the national headquarters staff itself.

This design contains a total of nine location domains and five management domains. To provide top-down management of the entire Unified Wireless Network, WCS Navigator is also deployed in the national headquarters location, providing both management and location visibility to all four regional campuses in addition to the national headquarters. WCS Navigator communicates with each of the five WCS servers via the WCS server's northbound API, allowing for the events, activities and resources of each regional management domain to be monitored and managed via a centralized portal. Figure 5-29 illustrates the network summary screen of WCS Navigator, where we can see its ability to monitor the status of multiple WCS servers as well as their location appliances. The alarm counts shown here represent aggregate quantities across all the monitored management domains. Clicking on any link takes the user to the appropriate detail screen using information retrieved from the WCS server that is responsible for the particular domain being queried.

Figure 5-29 WCS Navigator Summary Screen

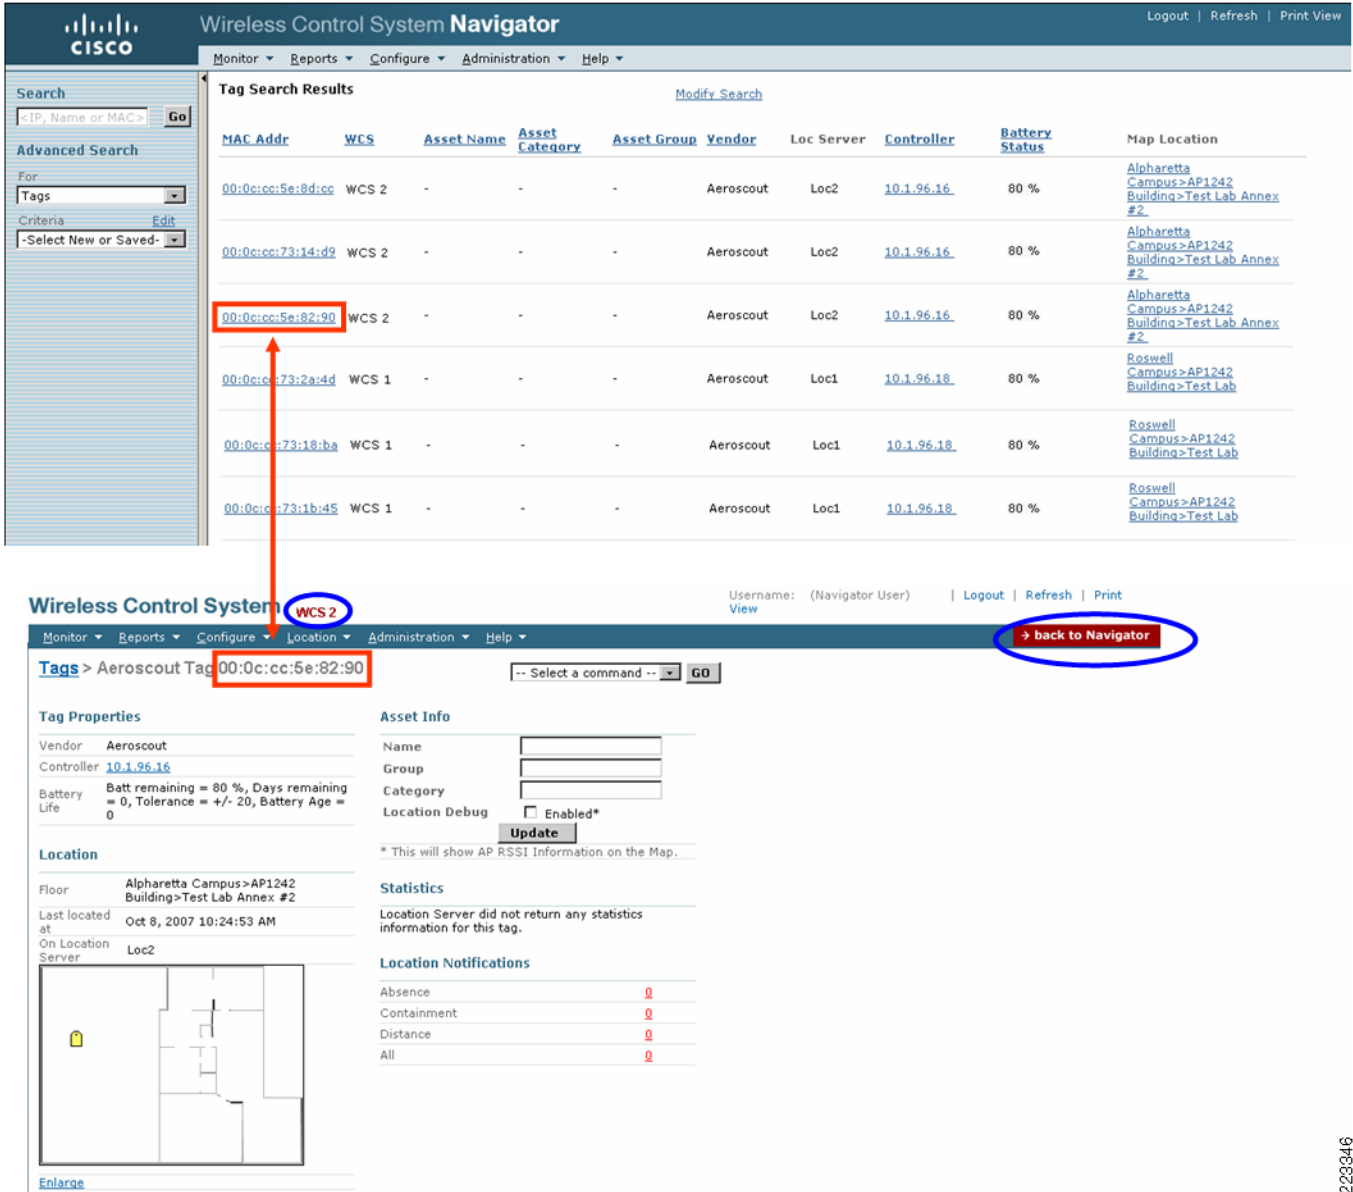


WCS Navigator allows for the location of individual clients, asset tags and rogues to be determined by searching across the entire set of WCS servers. Tracked device searches can be conducted based on IP address, user name, MAC address, or asset name, category or group. WCS displays the output of such searches in a list format (shown in the top half of Figure 5-30). Clicking on any of the tag MAC addresses, controller or map links transports the user to the appropriate detail screen using information from the responsible WCS server (as shown in the lower half of Figure 5-30). Note that when this occurs,

223945

additional fields (highlighted in blue ovals) are inserted into the WCS menus to indicate that the user has arrived at these WCS menus via the WCS Navigator and provide for an easy return path back to the main menu.

Figure 5-30 Result of Tag Search Across Multiple WCS Servers



The cautions stated earlier with regard to the degree and frequency of device migration between location domains also apply here. That is, attention should be paid to the level of tracked device migration that may occur intra-regionally, inter-regionally or between any of the regions and the national headquarters. In this case, when tracked devices migrate between location domains and are included in the active location database of two or more location appliances, duplicate entries may not only be seen in local WCS servers (as described earlier) but also in WCS Navigator. Figure 5-31 provides an illustrative

223346

example of this, for two tracked asset tags 00:0C:CC:73:14:D9 and 00:0C:CC:73:2A:4D. As described in the preceding section, judicious adjustment of the Absent Data Cleanup Interval can be used to reduce the lifetime of duplicate entries and help mitigate this condition.

Figure 5-31 Duplicate Devices in WCS Navigator Due to Device Migration

The screenshot shows the Cisco Wireless Control System Navigator interface. The 'Tag Search Results' table lists tracked devices. Two MAC addresses are highlighted: 00:0c:cc:73:14:d9 (yellow) and 00:0c:cc:73:2a:4d (red), indicating duplicate entries. The table columns include MAC Addr, WCS, Asset Name, Asset Category, Asset Group, Vendor, Loc Server, Controller, Battery Status, and Map Location.

MAC Addr	WCS	Asset Name	Asset Category	Asset Group	Vendor	Loc Server	Controller	Battery Status	Map Location
00:0c:cc:5e:8d:cc	WCS 2	-	-	-	Aeroscout	Loc2	10.1.96.16	80 %	Alpharetta Campus>AP1242 Building>Test Lab Annex #2
00:0c:cc:73:14:d9	WCS 2	-	-	-	Aeroscout	Loc2	10.1.96.16	80 %	Alpharetta Campus>AP1242 Building>Test Lab Annex #2
00:0c:cc:73:2a:4d	WCS 2	-	-	-	Aeroscout	Loc2	10.1.96.16	80 %	Alpharetta Campus>AP1242 Building>Test Lab Annex #2
00:0c:cc:5e:82:90	WCS 1	-	-	-	Aeroscout	Loc1	10.1.96.18	80 %	Roswell Campus>AP1242 Building>Test Lab
00:0c:cc:73:14:d9	WCS 1	-	-	-	Aeroscout	Loc1	10.1.96.18	80 %	Roswell Campus>AP1242 Building>Test Lab
00:0c:cc:73:18:ba	WCS 1	-	-	-	Aeroscout	Loc1	10.1.96.18	80 %	Roswell Campus>AP1242 Building>Test Lab
00:0c:cc:73:1b:45	WCS 1	-	-	-	Aeroscout	Loc1	10.1.96.18	80 %	Roswell Campus>AP1242 Building>Test Lab
00:0c:cc:73:2a:4d	WCS 1	-	-	-	Aeroscout	Loc1	10.1.96.18	80 %	Roswell Campus>AP1242 Building>Test Lab

223347

Antenna Considerations

Third-Party Antennas

When engineering in-building WLAN solutions, varying facility sizes, construction materials, and interior divisions can all pose concerns that need to be considered during design and deployment. Cisco Systems is committed to providing not only the best WLAN infrastructure and client components in the industry, but also providing complete WLAN solutions. To this end, Cisco Systems provides the widest range of antennas, cabling and accessories available from any wireless LAN manufacturer. With a full suite of directional and omni-directional antennas, low-loss cable, mounting hardware, and other accessories, installers and designers of Cisco-supplied wireless solutions that meets the requirements of some of the most challenging wireless LAN applications.

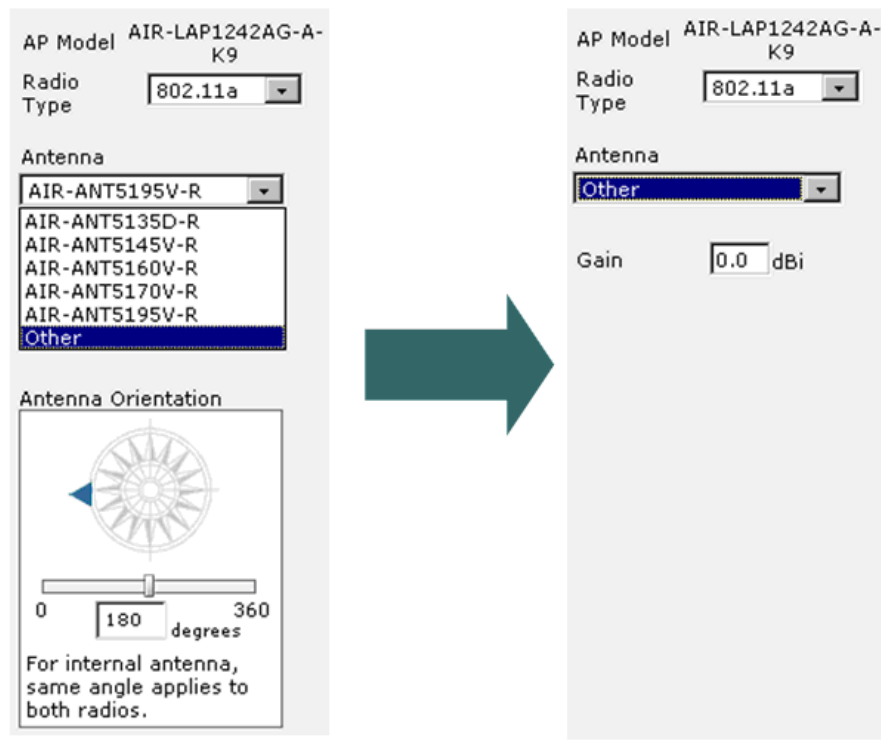
In the Cisco Unified Wireless Network, antennas available from Cisco Systems are pre-configured in WCS and available for assignment to access points via the drop-down menus found at Monitor > Maps > Position APs. Selecting a Cisco antenna from this list automatically defines the antenna's gain and propagation patterns to WCS and the location appliance, which helps facilitate optimal localization of tracked devices.

In some specialized cases however, there may be reasons to consider the use of third-party antennas that are not found on the WCS antenna list. These reasons may include:

- **Retrofit of a pre-existing installation**—If a pre-existing autonomous network is being upgraded to the LWAPP-based UWN solution, or if a pre-existing UWN installation is being upgraded, there already a large install base of third party antennas already deployed. Depending on their physical condition and their regulatory approval status for use with the latest 802.11 technologies, designers may wish to consider re-deployment.
- **Specific Product Requirements**—In some cases a specific physical or electrical requirement of the design might dictate the use of a niche third party antenna not contained on the WCS-supported list. For example, a fashion retailer may require the use of a “zero-footprint” antenna or an antenna available in a specific shape or color to augment the decor of a “Fifth Avenue” flagship retail location. Or an electronics manufacturing facility requires a directional antenna with a unique (and very specific) coverage pattern or polarization to better cover a specific area of the plant floor, while minimizing interference with sensitive equipment in a particular location.

WCS allows for antenna gain to be specified for antennas that are not on the dropdown list of standard antennas. This can be performed using the “Other” antenna option (shown in [Figure 5-32](#)). Custom azimuth and elevation propagation patterns for “Other” third-party antennas cannot be defined to either WCS or the location appliance, (note the loss of the antenna orientation compass in [Figure 5-32](#) when using the “Other” antenna option). Because of this, access points that are defined as being equipped with third party antennas will not be included in coverage heat maps and will not participate in client, tag, or rogue on-demand location tracking.

Figure 5-32 Specifying “Other” Antennas on WCS Floor Maps



190579

Since the exact propagation pattern of third-party antennas cannot be specified, a question often asked is whether the propagation patterns for a pre-defined Cisco Systems supplied antenna can be substituted when the antenna from Cisco Systems is regarded as being a “close match” for the third-party antenna. The answer to this popular question depends on several factors. For optimal location fidelity, it is recommended that one of the antennas listed within WCS be used whenever possible. However, if this is impossible, the following suggestions should be considered before performing such substitutions:

- For third-party antennas providing a gain of +6dBi gain or less, the difference in gain between the third-party antenna and the Cisco equivalent should not exceed 3 dBi.
- For third-party antennas providing greater than +6dBi gain, the difference in gain between the third-party antenna and the Cisco equivalent should not exceed 3 dBi. In addition, the gain of the third-party antenna must not exceed the gain of the Cisco equivalent. The latter condition must be enforced to avoid circumstances where excessive antenna gain may lead to regulatory compliance issues (FCC regulatory domain, other regulatory domains may differ).
- The third-party antenna should be of the same type as the Cisco equivalent. In other words, omni-directionals should only be substituted for omni-directionals, yagis for yagis, etc.

**Note**


Keep in mind that substitution of third-party antennas while configuring WCS for the nearest Cisco equivalent may not be supported by the Cisco Technical Assistance Center (TAC).

Antenna Orientation and Access Point Placement

When installing access points using either internal or external antennas, it is highly recommended that both the placement of the access point as well as the orientation selected for the access point antennas in WCS match the actual physical access point placement and antenna orientation. This helps to ensure accuracy and precision in both location tracking as well as the display of predictive heat maps.

The typical Cisco Aironet access point is installed using antenna diversity. Antenna diversity helps ensure optimal range and throughput in high multipath environments. With few exceptions, it is recommended that antenna diversity always be enabled. The location-aware Cisco UWN is designed to take RSSI information from both access point antennas into account when localizing tracked devices. For good accuracy, ensure that antennas are physically present on all enabled access point antenna ports. Failure to do so may cause inordinately low RSSI readings to be reported on enabled antenna ports that do not have an attached antenna. The use of abnormally low RSSI from antenna ports without antennas is not conducive to good location accuracy and should be avoided.

[Figure 5-33](#) illustrates how the configuration of the antenna's azimuth orientation within WCS is mapped to the actual physical orientation of the antenna. The blue triangle in the azimuth compass rose shown at the right of the figure indicates how the actual antenna should be physically positioned during deployment (notice that each of the antenna graphics contains a blue arrow as well). For omni-directional antennas, use unique identifying factors that are associated with the antenna (such as the right angled flexible antenna connector shown at the bottom of the 2.2dBi black whip antenna in [Figure 5-33](#)) to assist in proper positioning. For directional antennas, use unique physical characteristics of the antenna such as the exit location of the cable (for example, cable exiting up or cable exiting down) or other unique marks and construction characteristics.

Figure 5-33 Antenna Orientation

In software Release 4.1 of the location aware Cisco UWN, the ability to specify installed access point and antenna characteristics has been enhanced. Whereas prior releases assumed all access point antennas were uniformly installed at the same height on a floor, software Release 4.1 now provides the ability to account for antenna installations at varying heights. In this software release, the ceiling height that was configured when the floor was defined is used as the initial default for each access point. However, this can now be easily overridden on a per-access point basis (as shown in [Figure 5-34](#)). This capability allows the location appliance to incorporate varied access point antenna heights in its location calculations, an approach that more realistically approximates installations, especially in non-carpeted office type environments. Note that the individual height specified for an access point antennas cannot exceed the height of the floor.

Figure 5-34 Defining Individual Access Point Heights

Position access points on Floor Area 'Test Lab Annex #2'

	Horiz	Vert	AP Height	Zoom		
AP1242#4	42.3	24.7	12	100 %	Save	Cancel


Select each AP by clicking on it. Update its position, antenna information, height and when done with all APs click on Save.

**Note**

The antenna propagation characteristics of the AP1131 access point are optimal along its azimuth plane when ceiling mounted. For optimal location performance when using the AP1131, it is preferred that the access point be ceiling mounted rather than wall mounted.

In some cases, it is desirable to separate access points from antennas using a short length (less than 10 feet) of low-loss antenna cable. Reasons for this might include avoidance of obstacles or simply the desire to position access points and other active electronic infrastructure components within easy reach of local employees using commonly available ladders and stepladders. This facilitates easy removal and installation of these components should they require replacement. An example of this is shown in [Figure 5-35](#). In this case, a nationwide retailer has mandated that all electronic infrastructure components be accessible to store employees using the ten foot step-ladders commonly available at each store location. Here we see that the access point is mounted at 10 feet (for easy access) while the antennas are mounted at 15 feet. In cases such as this, the value specified for “AP Height” in [Figure 5-34](#) should reflect the height of the antennas and not the height of the access point.

Figure 5-35 Example of Differing Antenna and Access Point Heights



223850

Calibration

The Cisco WCS and the location appliance are shipped with default RF models that facilitate setup under two of the most common indoor office environments. One of these models represents a typical corporate office environment with both cubicles and drywall offices, and the other represents a similar environment with drywall offices only¹. These RF models provide an estimate of the path losses found in these typical indoor commercial office environments, and can be very useful when the primary requirement at hand is to provide a working location tracking system in the shortest amount of time possible.

Some indoor environments may possess more attenuation than is found in a typical office environment. In properly designed indoor installations where increased attenuation may be a factor contributing to less than optimal location accuracy, a site *calibration* may help restore lost performance. When an on-site calibration is performed, the system is allowed to sample path losses from known points throughout the environment, allowing it to formulate a custom RF model that provides a better understanding of the propagation characteristics specific to that environment.

1. A third RF-calibration model is provided for designers wishing to attempt outdoor deployments using the Cisco Wireless Location Appliance. This white paper does not address outdoor location design guidelines at this time.

In many cases, using the information collected during calibration instead of a default model can markedly reduce the degree of error between reported and observed client location. In environments where many floors share almost identical attenuation characteristics (such as the floors of a library containing similar arrangements of book shelves, for example), these strong similarities may allow for the RF model created by a calibration performed on any one of the floors to be applied to all floors with good results.

Calibration is actually a multi-step process that begins with the definition of a new calibration model via Monitor > Maps > RF Calibration Models > Create New Model. A step-by-step description of the calibration process can be found in the following two documents:

- “*Creating and Applying Calibration Models*” in the *Cisco Wireless Control System Configuration Guide*,
http://www.cisco.com/en/US/products/ps6305/products_configuration_guide_chapter09186a008083196d.html#wp1089489
- *Cisco Wireless Location Appliance: Deployment Guide* at the following URL:
http://www.cisco.com/en/US/products/ps6386/prod_technical_reference09186a008059ce31.html.

During the data point collection process, the calibration client repeatedly transmits probe requests on all channels. Depending on the particular calibration client being used, the client may be triggered to transmit probe requests on-demand via a network request. Clients that are incapable of recognizing these requests may be de-authenticated and disassociated in order to cause them to issue probe requests to the wireless medium and subsequently re-associate / re-authenticate. Access points in the vicinity of the client detect the RSSI of these probe requests and pass this information to their registered controllers. Controllers furnish the RSSI information detected during the calibration process to the location appliance and WCS for use in computing the path losses that will ultimately be associated with the new calibration model.

The data point collection phase of the calibration process in WCS can be performed using one of two methods. It can be performed from a single web-enabled mobile device associated to the WLAN, which controls both the probing of the network as well as the actual data collection. Alternatively, the data collection phase can be performed from two separate devices that are associated to the WLAN infrastructure. In this case, interaction with the WCS GUI is controlled from a primary device that is equipped with keyboard and mouse capabilities, while the actual generation of probe requests occurs on a second device.

Due to an open caveat¹ concerning the use of dual-band calibration clients and performing a location calibration data collection on both bands simultaneously, it is recommended that calibration data collection be performed for each band individually at this time. When using a dual-band client, use either of the following alternatives:

1. Perform the calibration data collection using a single laptop equipped with a Cisco Aironet 802.11a/b/g Wireless CardBus Adapter (AIR-CB21AG) on each band individually. For example, proceed to disable the 5 GHz band and complete the data collection using the 2.4 GHz band only. Then, disable the 2.4 GHz band and enable the 5 GHz band, and proceed to repeat the data collection using the 5 GHz band only.
2. Perform the calibration using two people and two laptops. Each laptop should have a Cisco AIR-CB21AG and be associated to the infrastructure using a different band. The two calibration operators may operate independently; there is no need for them to visit each data point together. In this way, a complete calibration data collection can be performed across both bands in half the amount of time as option #1 above.

1. For further information, refer to [CSCsh88795—CCX S36 Beacon Measurement Request Dual-Band Support, page 7-1](#).


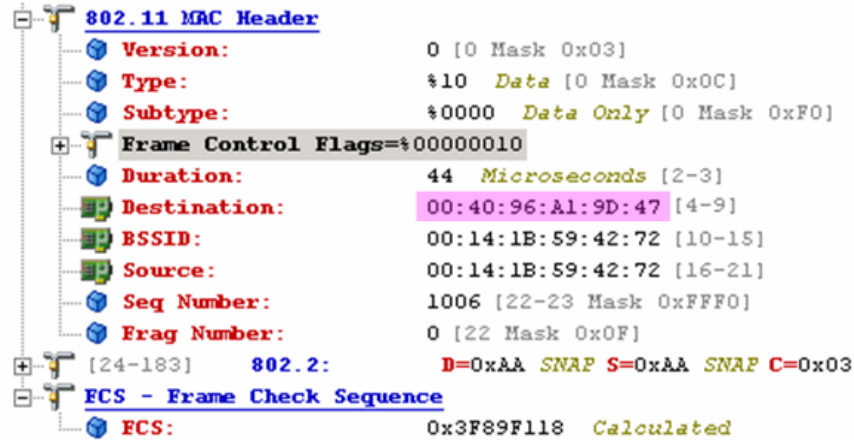
When a client compatible with the Cisco Compatible Extensions specification version 2 or greater is associated to the WLAN infrastructure and is specified as the calibration client in WCS, the client's MAC address will be inserted into the location calibration table of all controllers servicing the access points contained on the floor being calibrated. This insertion initially occurs immediately after the MAC address of the calibrating client and calibration campus, building and floor are specified via **WCS Maps > RF Calibration Models > Add Data Points**. After each save of a collected data point, the client MAC address will be removed from the controller's location calibration table. The client MAC address will then be briefly reinserted into controller location calibration tables upon each subsequent data point save and immediately removed thereafter. This process repeats for each data point collected.

When the MAC addresses of clients that are compatible with the Cisco Compatible Extensions specification version 2 or greater appear in the location calibration table of controllers, unicast Radio Measurement Requests will be sent to these clients (see [Figure 5-36](#)). Similar to how broadcast Radio Measurement Requests help improve the location accuracy of compatible clients during normal operation, unicast Radio Measurement Requests sent at short regular intervals (4 seconds) should cause compatible calibration clients to transmit probe requests very frequently. The use of Cisco Compatible Extensions Radio Measurement Requests and Cisco Compatible Extensions version 2 or greater clients allows this to occur without the need to force the client to continually disassociate and re-associate. This allows more consistent and reliable probing of the network, and allows smoother operation of the calibration client especially if it is being used as a workstation that is interacting with WCS via the calibration data collection GUI.

**Note**

Unicast Radio Measurement Requests will not be sent to clients that are associated to WLANs where the Aironet Information Element ("Aironet IE") has not been enabled on the supporting controller. For best calibration results with calibration clients supporting the Cisco Compatible Extensions specification version 2 or greater, ensure that Aironet IE support is enabled.

Figure 5-36 Unicast Radio Measurement Request



The fields in the unicast Radio Measurement Request shown in Figure 5-36 (highlighted with colored rectangles) are very similar to that of the broadcast Radio Measurement Request discussed previously in this document. However, Cisco Compatible Extensions Location Measurement broadcasts Radio Measurement Requests to all associated clients whereas during calibration data collection, Radio Measurement Requests are unicast only to calibration client MAC addresses that are contained in WLAN controller location calibration tables. For example, in Figure 5-36, we see that the MAC address we are unicasting the Radio Measurement Requests to is 00:40:96:A1:9D:47.

Calibration clients that are compliant with the Cisco Compatible Extensions specification version 1 (or not compliant with the Cisco Compatible Extensions specification at all) will not respond to unicast radio measurement requests. Instead, these clients will be forced to re-associate and re-authenticate in order to generate probe requests. If the Cisco Compatible Extensions Location Measurement parameter (discussed in the section entitled “Tracking Assets and Devices in the Cisco UWN”) is enabled on a controller for which a Cisco Compatible Extensions specification version 2 or greater client is being used for calibration, the calibration client should respond to both the broadcast request used for Cisco Compatible Extensions Location Measurement as well as the unicast Radio Measurement Requests used for calibration.



Note

The Cisco Aironet 802.11a/b/g Wireless CardBus Adapter (AIR-CB21AG) is recommended by Cisco Systems for location calibration data collection.

In most cases, WCS will handle the maintenance of the controller location calibration tables, inserting client MAC addresses as needed for calibration data collection and removing them when they are no longer necessary. Proper maintenance of this table is important, since every client compliant with the Cisco Compatible Extensions specification version 2 that is present in the table (up to a maximum of five) will receive unicast radio measurement requests whenever they are associated to any access point registered to that controller. This occurs regardless of whether the client is actively collecting calibration data points or being used simply as an associated mobile workstation during non-calibration periods.

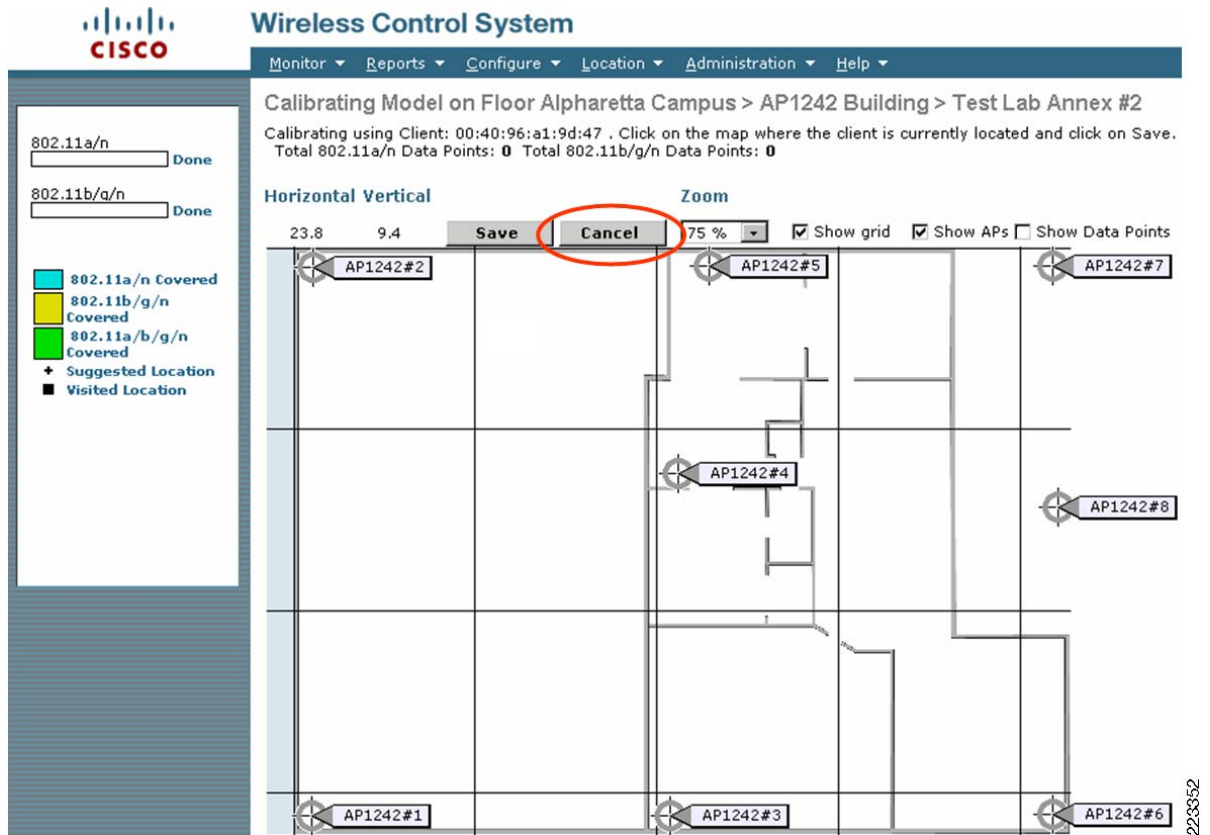
Because of this, residual client MAC addresses in controller location calibration tables can result in unnecessary load and traffic:

- Each WLAN controller containing these entries should perform the following every four seconds:
 - Issue radio measurement requests to each client in the table.
 - Process incoming probe requests from each client in the table.
 - Process incoming radio measurement responses from each client in the table.
- Each WLAN client present in the location calibration tables must perform the following every four seconds:
 - Issue one or more probe requests per channel in the regulatory channel set.
 - Process one or more incoming probe responses.
 - Generate a radio measurement response report containing the results of all probe responses received.

During normal calibration procedures, WCS will manage the addition and deletion of entries from the location calibration tables. However, there are some situations where residual stray entries may be observed. A common cause is improper termination of calibration data collection by the user, two examples of which might include:

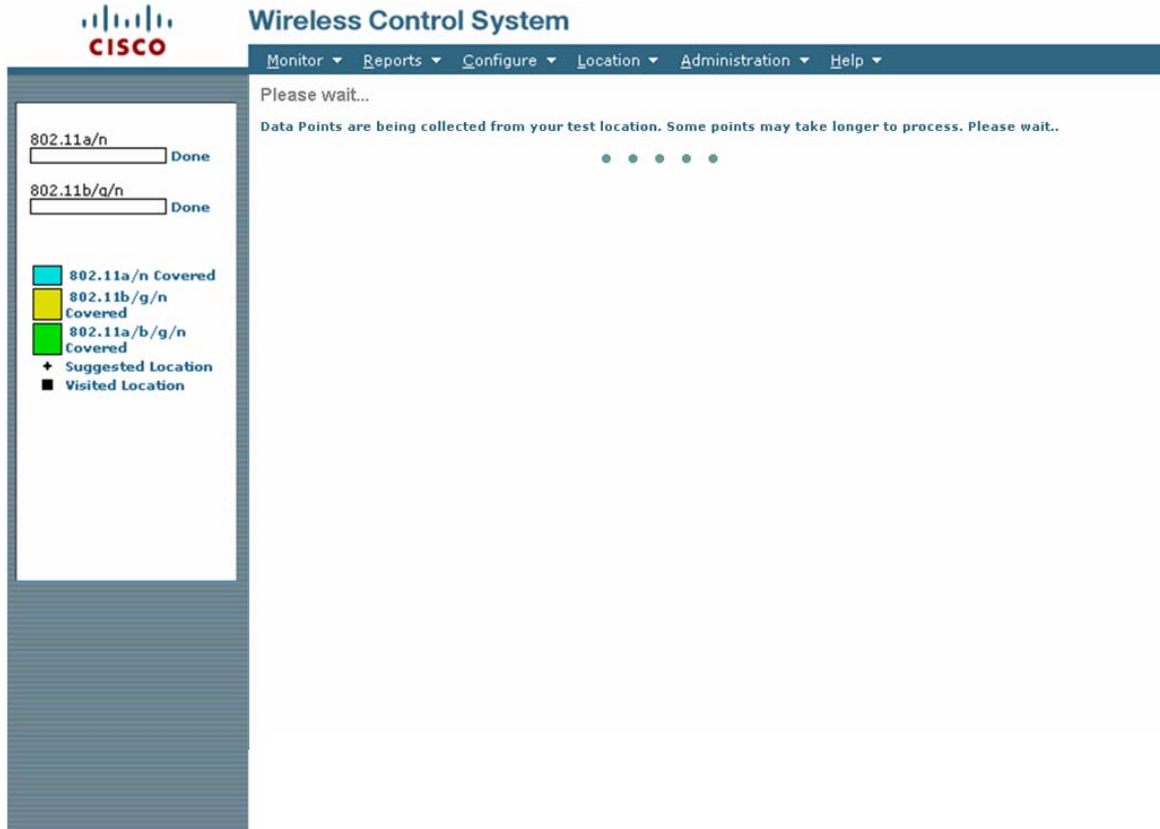
- WCS normally adds the client MAC address to the location calibration table immediately after the parameters on the Maps > RF Calibration Models > Add Data Points GUI menu panel are completed. If data point collection is aborted at this point by simply selecting a different WCS function from the main WCS menu, entries in controller location calibration tables will remain. To avoid this situation, it is recommended that users wishing to abort data collection use the “Cancel” option instead (shown in [Figure 5-37](#)) before returning to a previous menu.

Figure 5-37 Using the Cancel Option to Help Trigger Location Calibration Cleanup



- WCS adds the client address to the location calibration table immediately after clicking on the “Save” button shown in Figure 5-37 and initiating the data collection process for a particular data point. While data is being collected for the calibration client, the screen shown in Figure 5-38 will be displayed.

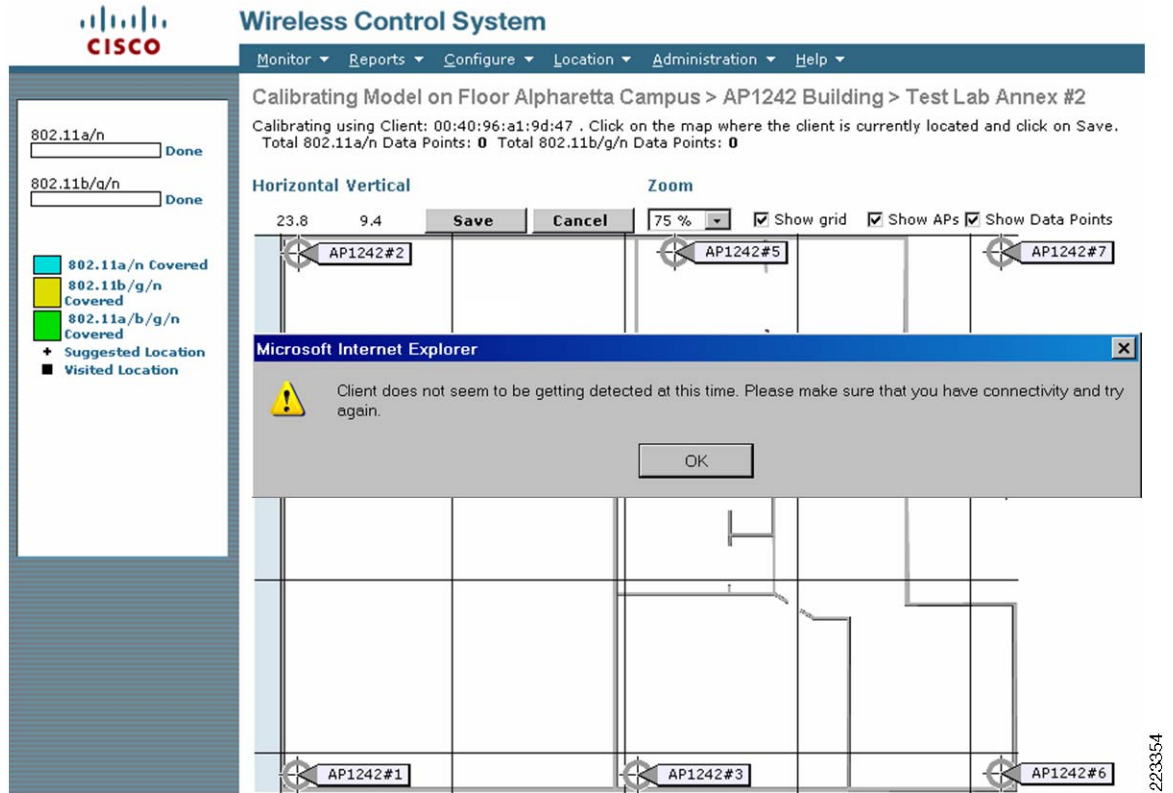
Figure 5-38 Data Point Collection in Progress



223959

Once the data collection has completed, the data collection screen will be re-displayed with the new data points illustrated on the map. In extraordinary situations, the data collection process may take a few minutes to complete (or timeout in about 90 to 120 seconds) due to an unexpected interruption in calibration client connectivity, such as moving into a coverage hole for example. If the user becomes impatient and decides to abruptly terminate the calibration by selecting a different menu option from the WCS main menu, residual controller location calibration table entries may occur. To avoid this, it is recommended that all calibration users refrain from this behavior and instead exercise patience in waiting for the data collection to complete or timeout (shown in [Figure 5-39](#)).

Figure 5-39 Data Collection Timeout Message



In cases where the contents of the controller's location calibration table is in question, it can be queried directly using the CLI command.

```
(Cisco Controller)>show client location-calibration summary
  MAC Address      interval
-----
00:40:96:a1:9d:47    4
```

In situations where there are unnecessary residual entries present in a controller's location calibration table from prior calibration data collections, they can be removed manually using the following controller CLI command:

```
(Cisco Controller)>config client location-calibration disable
00:40:96:a1:9d:47
```


**Note**

Before removing any entries in controller location calibration tables, it is good practice to be sure that these entries are not in use by other users performing a calibration data collection.

Calibration data collection should be performed after the system has been fully installed, basic coverage checks are completed, and the recommended RSSI cutoff (typically -75 dB or better) to a minimum of three access points has been verified. All access points should be in place, properly oriented and registered to their respective controllers with WCS and the location appliance fully operational. For optimum ease of use and visibility during the data collection procedure, a portable laptop or tablet computer (or “computer on wheels” cart setup) with a large, clear and bright screen is recommended, especially in areas of bright ambient light.

During the calibration data collection process, WCS suggests locations on the floor map where samples should be taken (shown in Figure 5-40) along with graphical indication of the degree of progress achieved. At any point during the calibration data collection process, a graphical representation of the calibration points captured thus far can be generated by clicking on the dropdown menu in the upper right hand corner of the calibration data collection screen. The calibration data collection process can be completed in one session, or the session can be stopped and returned to at a later time. This process can be repeated as often as necessary to complete calibration data collection for a floor.

Figure 5-40 Calibration Data Collection Screen



223365

Calibration Validity

Strictly speaking, a properly performed site calibration is considered valid as long as the fundamental environmental factors affecting RF propagation between clients and access points have not deviated significantly from the state under which the original calibration was performed. For example, significant changes in the material contents of the target environment may have an impact on the path losses experienced within that environment. Performing a re-calibration allows the system to better understand the current level of attenuation and fading present in the environment and allow it to re-calculate the path loss model. In many cases, this can help to restore lost accuracy and performance to the system.

From a practical standpoint, the location-aware Cisco UWN is seen to be more tolerant of environmental path loss changes when those changes move from higher path losses to lower path losses, relative to the path loss model currently in use. Therefore, if faced with designing a location tracking for an environment where seasonal content variations are the norm, it is recommended that the calibration of the site be performed at peak content levels, when path losses would be expected to be at their highest. The use of a high path loss model when actual path loss is lower than expected has been shown to produce better location performance than the use of a path loss model with low path loss when the actual path loss of the environment is higher than expected.

Examples of the types of changes where a re-calibration may be recommended if a significant drop in performance is noticed include, but are not limited to, the following cases:

- Changes in stocked material—A floor of an supply warehouse that was last calibrated when it contained paper products has now been converted to stocking bulk metal containers of dense liquids, such as motor oil, transmission fluid, gear oil, etc).
- Changes in interior walls—A newly remodeled office facility was last calibrated prior to the installation of several new walls with improved fire protection and interior sound deadening insulation.
- Changes in stocking density—A large library was originally calibrated when it was still using older bookshelves that contained six to eight shelves per stack. However, it has since been upgraded and now sports new bookshelves that contain between ten to twelve shelves per stack.
- Changes in access point density—A manufacturing site was originally calibrated for location tracking with 50 access points and an inter-access point spacing of 40 feet. However, due to a business slowdown, a large portion of the plant has been mothballed with 50% of the access points powered down. The effective inter-access point spacing at this point is 65 feet.

Tips for Successful Calibrations

Data Collection

As stated earlier, the WCS calibration process helps ensure that a sufficient number of calibration data point measurements are collected before allowing the calibration user to move forward with calibrating the model and applying it to floors. During the calibration process, use the blue crosshairs on the calibration grid (shown in [Figure 5-41](#)) as suggestions with regard to where the calibration client should be positioned when collecting data points. Always make sure you accurately position the red cross hairs prior to clicking on **Save** and initiating data collection.

Although they are only suggested locations, the blue crosshairs are an excellent way to stay on track and uniformly cover ground, especially within large environments. The calibration grid will be updated to indicate the locations actually visited, with the surrounding area of localization that is now “covered” being indicated by a blue color for 802.11b/g, yellow for 802.11a, and green for both bands, as illustrated in [Figure 5-41](#).



Note

In order to promote better location fidelity, every attempt should be made to be as accurate as possible when indicating the calibration client's actual physical position using the red crosshairs in [Figure 5-40](#) and [Figure 5-41](#).

Figure 5-41 Example of a Completed Calibration Data Collection for Both Bands

Calibration Status

802.11a Done

802.11b/g Done

Legend

- 802.11a Covered
- 802.11b/g Covered
- 802.11a,b/g Covered
- + Suggested Location
- Visited Location

Rogues

Coverage	0	19	0
Security	2	0	0
Controllers	2	0	0
Access Points	18	0	0
Location	0	0	0

Calibrating Model * Berlinger Suburban Office

Building > Test Lab Annex #2

Calibrating using Client: 00:0f:b5:10:62:97 . Click on the map where the client is currently located and click on Save.
Total 802.11a Data Points: 258 Total 802.11b/g Data Points: 259

Horizontal Vertical Zoom

0 0 Save Cancel 75 % Show grid Show APs Show Data Points

Keep in mind that the calibration utility is not able to recognize floor plan obstructions or hazards, such as interior walls, pipes, racks, or other structures. Therefore, it is not unusual to have a suggested data point crosshair appear in an area that is physically inaccessible to the operator. In that case, simply visit a location as close as possible to the inaccessible location and perform the calibration data collection there. Make sure, however, that the red cross hair is positioned to correctly indicate the physical location where the data collection actually took place.

Calibrating Under Representative Conditions

As mentioned previously, the location appliance and the Cisco WCS use the information gathered during a calibration data collection to better understand the propagation characteristics present within the environment. This information is culled from the aggregate of all the data points accumulated during calibration data collection. To facilitate an accurate calibration, it is recommended that the environment in which the calibration data collection is performed be representative of the daily production environment. For environments that experience variations in the level of material and personnel present, it is recommended that calibration be performed at a time when such levels are at or near their peak.

For example, the calibration should be performed during business hours when the facility contains a representative population of people (human attenuation) as well as material on shelves (material attenuation). If carts, racks, beds, or other large metallic objects are normally used in this environment, these should also be present during calibration. If large doors are present in the environment, they should be positioned as they would normally be during business hours. In most cases, calibrations performed during normal business hours are more likely to be representative of the levels of attenuation found in the production daily environment than off-hours calibration.

Often, the most convenient time to perform a calibration may be after construction is completed but before people and contents are moved into a site. Cisco Systems highly recommends against performing calibrations of such “empty rooms”. Once these areas are stocked and occupied, attempting to localize tracked devices using a RF model based on data collected in an empty and barren environment is not likely to provide optimal results. When presented with a choice between calibrating when stockroom shelves are only at half capacity or at full capacity, the calibration that is done at full capacity will typically yield better accuracy, even when used at times when the stockroom is only half full.

If it is necessary to deliver an operational location tracking system on “day one” for a newly constructed area, you may wish to use one of the RF models that are supplied with WCS and the location appliance, as a temporary measure. Once the area has been fully stocked and staffed, perform calibration data collection under conditions that would be considered representative of its peak or normal capacity. After the calibration has been completed, use the newly created RF model instead of the supplied model chosen originally. In this way, the supplied model initially chosen allows for users and administrators to quickly familiarize themselves with the system, and the subsequent switch to a properly calibrated RF model should provide for better overall performance.

In order to plan for the most optimal time to perform a calibration of the area, the designer should work closely with those personnel possessing an intimate knowledge of the business patterns and processes occurring there. This is especially true if seasonal variation in stocking levels may occur, as would be the case in a retail or logistics site. If it cannot be determined from prior conversations with site personnel, one way to determine a good time to perform a calibration is to visit the site beforehand and observe the activity pattern of both the facility and the personnel present. Prior observation of activity in this manner allows the designer to plan for the optimum time to perform the calibration, so as to yield the most representative results and also to not excessively inconvenience the personnel employed at the facility.

Recommended Calibration Clients and Techniques

The Cisco Aironet 802.11a/b/g Wireless CardBus Adapter (AIR-CB21AG) is highly recommended for use as a calibration client by Cisco. This client is compatible with the Cisco Compatible Extensions specification for WLAN devices at version 5, and fully supports the use of both broadcast and unicast radio measurement requests as described earlier in this section.

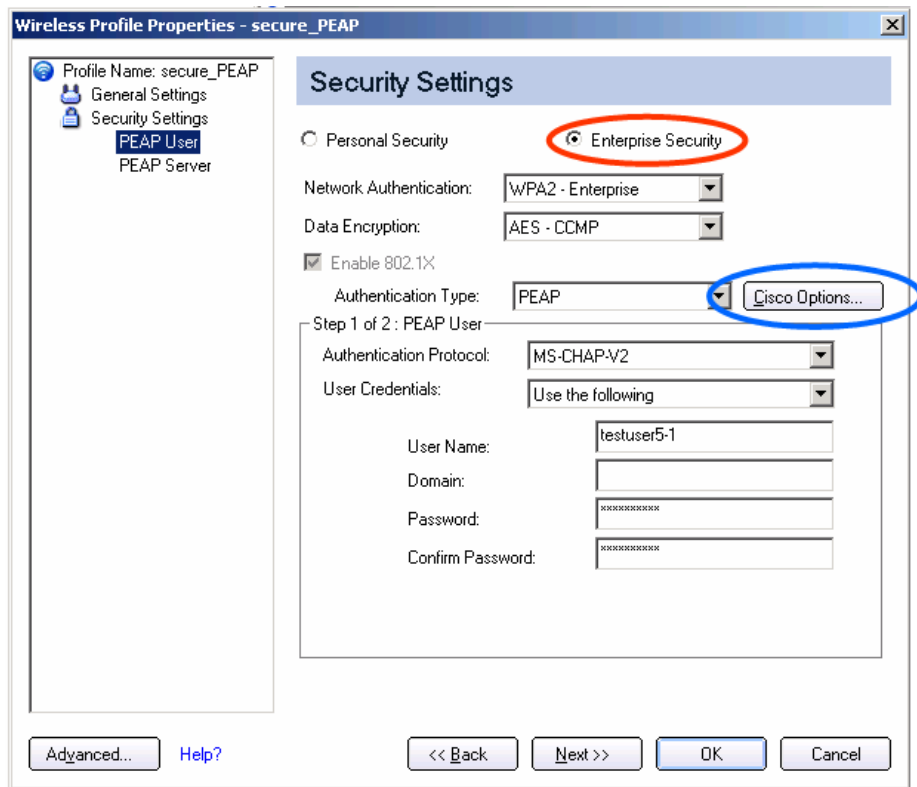
If a Cisco Aironet 802.11a/b/g Wireless CardBus Adapter (AIR-CB21AG) cannot be used as a calibration client, a third party client device that is compatible with the Cisco Compatible Extensions specification for WLAN devices at version 2 or higher may be used. In order to assure a reliable calibration data collection, any third party WLAN client used for location data collection should be capable of recognizing and responding to S36 unicast radio measurement requests sent during calibration. Also, the transmit power level used by third party clients when transmitting probe requests should be known.

WLAN client devices that are not compatible with the Cisco Compatible Extensions for WLAN devices specification, or compatible only with version 1 of the specification, are not considered optimal for use in location calibration data collection.

Note that when using a laptop computer containing the Intel® PRO/Wireless 3945ABG Network Connection or the Intel® PRO/Wireless 2915ABG Network Connection adapter, the default configuration is for a *Personal* level of security settings (intended for non-enterprise use) that does not include compatibility with the Cisco Compatible Extensions specification. When using this default *Personal* level of wireless security, clients equipped with the Intel 3945ABG or 2915ABG client adapters will not support S36 unicast or broadcast radio measurement requests and are not compliant with the Cisco Compatible Extensions specification for WLAN devices.

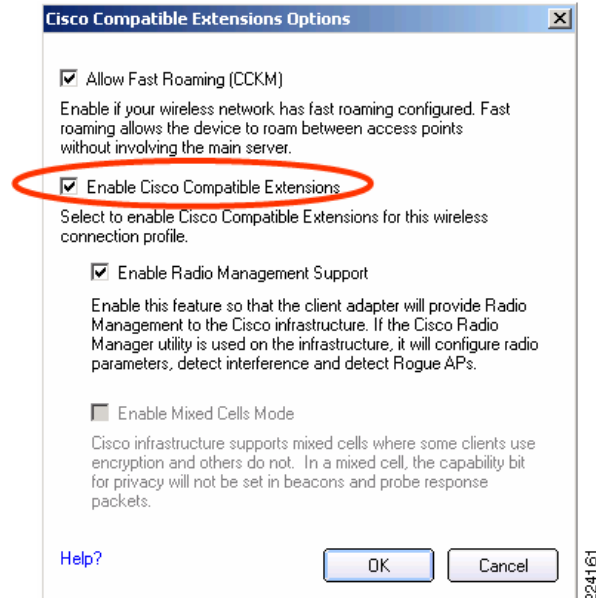
In order to enable compatibility with the Cisco Compatible Extensions specification, the Intel ProSet client supplicant must be used to reconfigure the client for *Enterprise Security* and enable Cisco Compatible Extensions. [Figure 5-42](#) and [Figure 5-43](#) illustrate how this is performed.

Figure 5-42 Intel ProSet Security Settings Panel



In [Figure 5-42](#), under the Security Settings panel of the Intel ProSet configuration, select the **Enterprise Security** option (highlighted by the red circle) instead of the default **Personal Security** option. Configure the appropriate authentication and authentication types, and select **Cisco Options** (highlighted by the blue circle). This will present the panel shown in [Figure 5-43](#).

Figure 5-43 Intel ProSet Cisco Compatible Extensions



In Figure 5-43, Cisco Compatible Extensions should be enabled (as indicated by the red circle). Note that Cisco Compatible Extensions is automatically enabled when configuring profiles for CKIP, LEAP, or EAP-FAST.

**Note**

For additional information regarding the configuration of the Intel® PRO/Wireless 3945ABG Network Connection or Intel® PRO/Wireless 2915ABG Network Connection adapters, refer to the following documents from Intel Corporation: <ftp://download.intel.com/support/wireless/wlan/sb/3945abgug.pdf> ftp://download.intel.com/support/wireless/wlan/pro2915abg/sb/2915ABG_UG.pdf

Calibration should be performed using a calibration client and a suitable laptop computer with a fully charged battery. The following recommendations should be considered when performing calibration data collection:

1. Ensure that your calibration client is being detected by the access points on the floor where you wish to perform the calibration.
2. Temporarily disable Dynamic Transmit Power Control (DTPC) prior to conducting calibration data collection. DTPC must be disabled separately for each band using either the controller GUI, the controller CLI or WCS for each controller whose registered access points are expected to participate in calibration data collection. After calibration data collection has been performed, DTPC should be re-enabled for normal production operation.
3. Ensure that the WLAN to which your calibration client will associate is configured to support Aironet Information Elements (Aironet IE). Doing so will enable the use of unicast radio resource measurement requests during calibration data collection for more efficient operation.

To obtain best performance when displaying access point coverage heat maps and tracking devices in most cases, calibration clients should be pre-configured as closely as possible to transmit power levels of 63mW (+18dBm) for 2.4GHz and 32mW (+15dBm) for 5 GHz. Note that it is imperative that DTPC be disabled (item number two above) such that these transmit power levels do not vary during calibration data collection.

Due to an open caveat¹ concerning dual-band calibration clients when attempting to perform a simultaneous data collection on both bands, it is recommended that calibration data collection be performed for each band individually at this time. In order to do this using a dual-band client, use either of the following alternatives:

1. Perform the calibration data collection on each band individually using a single laptop equipped with a dual-band client adapter compatible with the Cisco Compatible Extensions specification for WLAN devices specification at version 2 or higher, and capable of recognizing and responding to S36 unicast radio measurement requests. An example of such a client is the Cisco AIR-CB21AG. For example, proceed to disable the 5 GHz band and complete the data collection using the 2.4 GHz band only. Then, disable the 2.4 GHz band and enable the 5 GHz band, and proceed to repeat the data collection using the 5 GHz band only.
2. Perform the calibration data collection using two operators and two independent laptops. Each laptop should be equipped with a dual-band client adapter compatible with the Cisco Compatible Extensions specification for WLAN devices specification at version 2 or higher, and capable of recognizing and responding to S36 unicast radio measurement requests. An example of such a client is the Cisco AIR-CB21AG. Each laptop should be associated to the infrastructure using a different band. The two calibration data collection operators may function independently, there is no need for them to visit each data point at the same time, or to even visit the same data points. In this way, a complete calibration data collection can be performed across both bands in half the amount of time compared to option #1 above.

Some embedded laptop client adapters may not transmit probe requests at these power levels, but instead are restricted to lower transmit power levels (for example, +15dBm for 2.4 GHz). Generally, these clients can still be localized with acceptable accuracy when calibration data collection is performed according to the guidelines outlined above. However, a slight increase in location accuracy may be possible if the calibration data collection is performed at power levels with which we expect the embedded laptop client to transmit its probe requests. A tradeoff that must be considered when opting for this approach, is the possibility of reduced access point heat map accuracy (heat maps are most accurate when calibration is performed at 63mW (+18dBm) for 2.4GHz and 32mW (+15dBm) for 5 GHz).

After calibration data collection has been completed, all temporarily configured parameter changes should be returned to their normal settings. In some cases, the device that is used to control the calibration data collection process may not be the same device that is used to transmit probe requests. For example, a laptop with an embedded wireless adapter compatible with the Cisco Compatible Extensions specification version 1 might be a user's preferred device based on an ergonomic or special accommodation feature that he or she requires. Since this device is not compatible with the Cisco Compatible Extensions specification at version 2 or greater, its use is not recommended for optimal results when performing calibration data collection. However, we can use this device to log into the UWN and control the data collection process remotely, assigning the role of transmitting probe requests to another device that is, for example, equipped with a Cisco AIR-CB21AG. As shown in [Figure 5-44](#), this arrangement allows the location of the probing client (otherwise referred to as a “remote calibration client”) as well as the timing of each data collection to be fully controlled from the laptop. The probing client (with MAC address 00:1a:a1:92:a1:20 in [Figure 5-44](#)) is remotely instructed to issue probe requests to the network infrastructure appropriately.

1. For further information, refer to [CSCsh88795—CCX S36 Beacon Measurement Request Dual-Band Support, page 7-1](#).

Figure 5-44 Controlling Calibration Data Collection Remotely

The screenshot shows the Cisco Wireless Control System (WCS) interface. The top navigation bar includes 'Monitor', 'Reports', 'Configure', 'Location', 'Administration', and 'Help'. The main content area is titled 'Calibration Model > 'Test #3' > Start Calibrating'. It features a form with the following fields:

- Enter MAC Address of Client* ***: A text input field containing '00:1a:a1:92:a1:20'.
- Choose the Floor on which this Model is intended to be calibrated**: A section with three dropdown menus:
 - Campus**: 'Alpharetta Campus'
 - Building**: 'Concannon Building'
 - Floor Area**: 'Main Floor'

Below the form are 'Next' and 'Cancel' buttons. A note at the bottom states: '* Client should be detected by APs on the chosen floor. For calibration, Automatic power assignment should be turned off. This can be done by making sure that Tx Power assignment mode for the Radios(802.11a/n & 802.11b/g/n) on the selected floor is set to Custom OR the controllers' Dynamic Power Assignment is set to Disable. After you are done with calibration, you can turn on the automatic power assignment.'

On the left side of the interface, there is a sidebar with a 'Quick Search' section containing a search box and a 'Go' button, and a 'Search Maps' section with a 'New Search...' button and a 'Saved Searches' dropdown menu.

223356

Calibration of Non-Uniform Environments

In some cases the network designer is faced with challenges because of an environment that is of non-uniform construction. An example is a single floor consisting of a large call center cubicle area (path loss exponent of 3.3), dense metal racking and electronic equipment in a second area (path loss exponent of 4.3) and a large group of individual offices with drywall walls in a third area (path loss exponent of 3.5).

Cases such as this can be addressed via one of two options:

1. Calibrate in the area with highest expected attenuation (path loss)—The most straightforward method in which to handle this situation is to perform the calibration in the areas possessing the highest overall attenuation (i.e. the highest path loss exponent), and apply the resulting RF model to all areas of the floor. In mathematical simulations as well as lab research, the application of an RF model that is based upon a higher level of path loss to areas where the actual path loss is lower has shown to provide better location accuracy than the converse approach. Thus in our example, it would be recommended that the calibration be performed in the area with the dense metal racking and electronic equipment (path loss exponent of 4.3) and the RF model that results from this calibration used for the entire floor.
2. Calibrate across all areas of the floor—This approach takes into account all areas of the floor and attempts to produce a “balanced” RF calibration model. While performance may be acceptable using this approach depending on the accuracy needs of the location application, laboratory testing indicates that in general, improved results are obtained when using option one above. Note that if there are large differences in size between the different floor areas that result in significant differences in the number of calibration data points collected within each area, the final path loss model using this approach may be biased in the direction of the path loss associated with the larger areas.

3. Address the different areas of the floor as if they were individual floors—In some cases, improved accuracy can be obtained within each individual area with a tradeoff of increased management overhead and some potential edge accuracy degradation. Since WCS and the location appliance do not allow for the provisioning of different RF models to sub-floor areas in software Release 4.1, each of these sub-floor areas would need to be defined as a separate floor in WCS. Individual calibrations are then performed in each of the sub-floor areas, applied to their pseudo-floor definition within WCS and then synchronized with the location appliance. Because this approach allows for the provisioning of separate path loss models that are attuned to the characteristics of each sub-floor area, improved accuracy and precision is possible. However, potential tradeoffs may include additional management overhead on the part of the WCS administrator. An organized naming convention is typically required for floors and sub-floors such that they are easily recognizable by WCS users and able to be logically considered as a group. Each sub-floor area should be considered as an independent location area subject to the location-aware design recommendations made in this document. Also, it should be noted that accuracy may degrade as devices approach the edges and borders of sub-floors, since the location appliance positioning engine does not consider signal strength readings from access points that are resident on a different floor in software Release 4.1.

In addition to these two mainline options, it is possible to treat each of the three areas as separate “floors” in WCS, and thereby allow the development of RF models attuned to each area's characteristics. However, this approach possesses a serious limitation that must be understood. In addition to the management overhead of developing floor and sub-floor naming conventions that bring a modicum of sensibility to such an approach, a technical limitation exists whereupon the location appliance currently does not take into account tracked device RSSI coming from access points that are deemed to be located on a different floor than the tracked device itself. Thus, when physical floors are divided into areas that are then defined to WCS as individual floors themselves, tracked devices that venture into the edge boundary areas of these newly defined “floors” may experience degraded accuracy. Unless the designer as well as the system user is comfortable with this limitation and its potential impact on boundary area accuracy, this approach is best avoided.

While there is no *ideal* solution to situations where the degree of uniformity is vastly different across a floor, in general option #1 is observed to offer the best compromise between ease of implementation and performance.

Inspecting Location Quality

Location inspection allows path loss model accuracy to be validated by comparing the actual versus the predicted location of calibration data points. Unlike the location planner or location readiness tools, which are purely predictive in nature, when you perform location inspection, you are directly comparing the predicted locations of calibration data points to the actual physical locations originally specified by the calibration operator. In order for location inspection to deliver on its true value, however, accurate placement of the red crosshairs during calibration data collection is very important.

The results of each data point comparison are used to graphically express the overall accuracy of the path loss model at various points on the floor. This provides the system designer or installer with “real-world” feedback with regard to how the expected performance of the system compares to its actual performance, given the calibration client used and the condition of the environment at the time of calibration.

Location inspection is accessible from the **Monitor > Maps > RF Calibration Model > model name** WCS menu via the “Inspect Location Quality” hyperlink located next to the name of the floor where data collection for the calibration model was performed, as shown in [Figure 5-45](#).

Figure 5-45 *Accessing Location Inspection*

[Calibration Model](#) > Beringer Suburban Office 2

Status	Calibrated
Last Calibrated On	July 5, 2006 11:18:19 PM EDT
Total 802.11a Data Points	258
Total 802.11b/g Data Points	359
802.11a Calibration Done	Yes (154 % done)
802.11b/g Calibration Done	Yes (152 % done)

Calibration Floors
 Alpharetta Campus > AP1242 Building > Test Lab Annex #2 ([Inspect Location Quality](#))

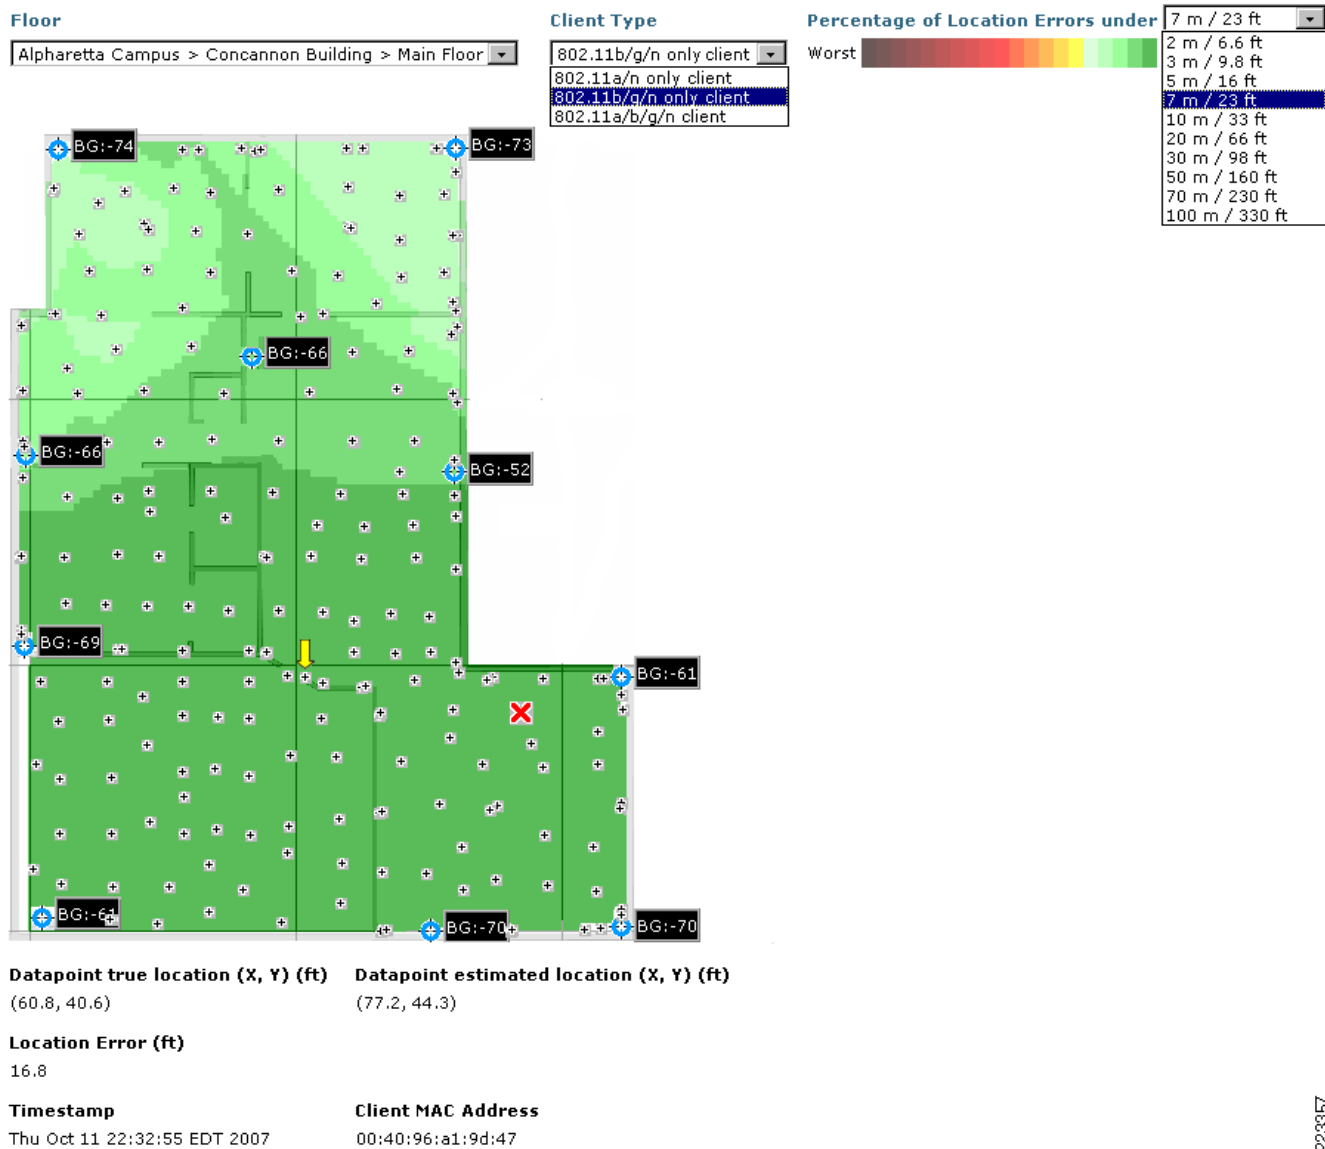
Floors Applied To
 Alpharetta Campus > Building 935 > Atrium Floor
 Alpharetta Campus > AP1242 Building > Test Lab Annex #2

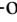
190584


Figure 5-46 illustrates an usage example for the Location Inspection tool. Here we have performed a location inspection after a calibration has been completed for a test lab facility. The results indicate the level of accuracy and precision the location appliance delivered during the calibration. In this example, using an 802.11bg-only calibration client, the location appliance is seen as capable of delivering a level of accuracy and precision (based on conditions in place at the time of calibration) of 7 meters or 23 feet with 90% precision over the majority of the test lab area. Looking at a single calibration data point in specific (artificially indicated in the figure by a yellow arrow to represent the point at which a mouse-over was performed), we see the estimated location indicated by **X**, as well as the details behind the degree of location error for this particular case (16.8 feet).

Figure 5-46 802.11bg Location Inspection

Inspect Location Quality > Calibration Model > Concannon Test Lab

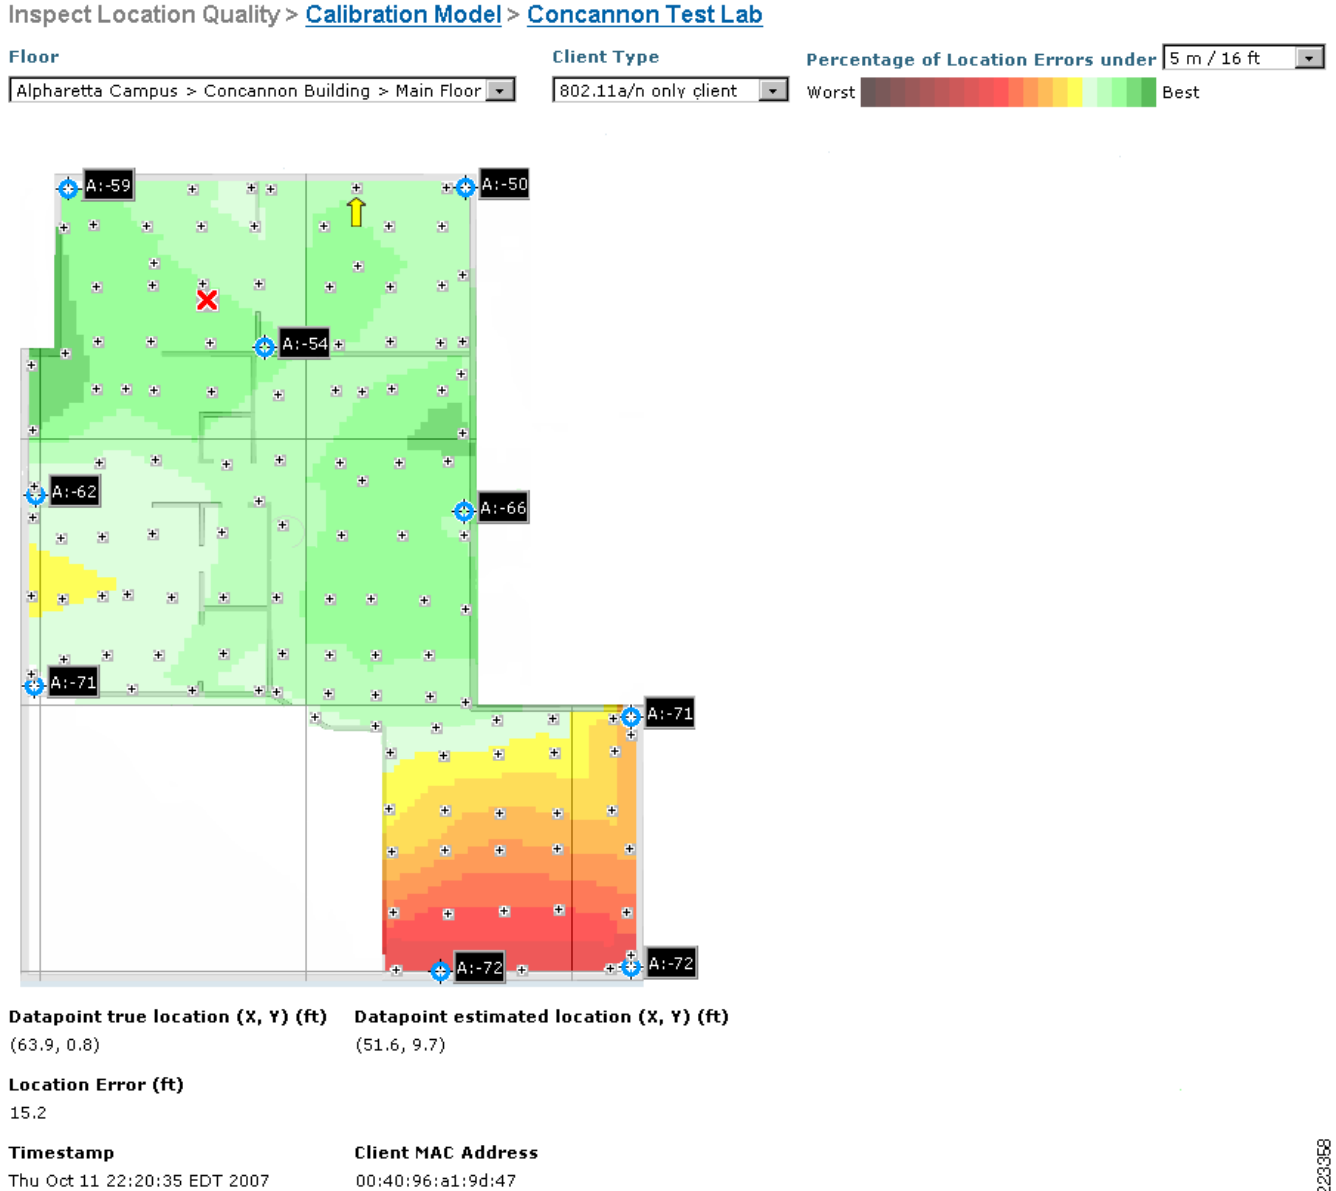


When a mouse-over is performed of any calibration data point (whose location is indicated using small crosshair  and is based on the actual coordinates reported by the calibration operator), the following information is displayed:

- The predicted location of the calibration client, depicted on the location inspection display by 
- The RSSI detected by each contributing access point and the band the contribution was made on, as shown in Figure 5-46 by symbols such as BG=-74 or A=-59.
- The true and estimated location of the calibration data point listed numerically, in terms of x and y coordinates.
- The estimated location error.
- A time stamp indicating when the calibration information was collected and the MAC address of the calibration client used to perform the collection.

Note that the band can be specified (2.4 GHz, 5GHz, or both) as well as the performance criteria. The flexibility imparted by this level of control allows the system designer to perform “what if” planning based on the results of location inspection, and examine the limits of higher (or lower) levels of accuracy and precision. This can be useful, for example, when planning for future location applications or in analyzing what areas of the current environment might require additional attention. For example, it is quite easy for us to use both of these controls to visualize the limits of location precision at the 5 meter accuracy level using 802.11a instead of 802.11bg, as shown in Figure 5-47.

Figure 5-47 802.11a Location Inspection Example at 5 m Accuracy



In the case of our test lab example in Figure 5-47, we see that the infrastructure appears to be capable of delivering 5 meter accuracy using 802.11a, with a precision of 85% percent or better in the top two thirds of the test lab floor. However, the bottom area of the figure reveals that challenges exist in meeting this level of accuracy in the lower third of the floor, thereby meriting further investigation. The value of

location inspection here is that the modeling of future scenarios where increased accuracy may be required can be performed using the information collected during a current calibration, without involving actual users and without requiring them to participate in trial and error testing.

In both [Figure 5-46](#) and [Figure 5-47](#), note the appearance of calibration data points all along the perimeter of the test lab, between the perimeter access points. These data points were purposely placed there in order to eliminate (or at least reduce) the appearance of “white space” along the perimeter during location inspection. When pure white spaces occur in location inspection output, often times it is due to a lack of data in those areas, which prevents location inspection tool from calculating a valid representation of accuracy and precision at those points. Rather than attempt to provide an estimation based on little or no information, the location inspection tool leaves these areas blank.

To help avoid such behavior, it is suggested that calibration data points be taken along the perimeter, as well as in areas contained within the perimeter. Perimeter data points can be added after the initial calibration if desired. To do this, simply rerun the “Add Data Points” data collection phase for the calibration model, and be sure to take a sufficient number of new data points directly in these white areas. After completion, rerun the “Calibration” phase and re-inspect location quality. These white areas should now be totally eliminated or at least significantly reduced. The process can be repeated if necessary to further address any remaining white areas if still present.


Note

Signal strength information for each selected band as well as the test client’s actual location coordinates must be available for each floor targeted by the location inspection tool. Even if multiple floors share the same RF model, only the floor upon which the model was actually calibrated is eligible for location inspection.

Using Test Points to Verify Accuracy

Complementing the capabilities found in Location Inspection, beginning with software Release 4.1 the location aware Cisco UWN allows for impromptu “go/no-go” verification of whether the location appliance’s baseline accuracy and precision specification has been met on a particular floor. This can be done using the *test point facility*, which is accessible from the WCS main menu via **Monitor > Maps > floor map name > Position Test Point**. For a specified device MAC address, the test point facility keeps track of the total number of location test point samples taken, and can calculate the percentage of the total number of test point samples taken whose location accuracy are within the location appliance’s baseline specification.

The test point facility is useful when there are tagged assets, wireless client devices or even rogues on a floor whose actual physical location is known and which move very infrequently or not at all. An example of this might be asset tags that are deployed attached to shipping containers that will not be used during this shift, or a wireless-equipped desktop computer that does not move from its deployed location at a supervisor’s workstation. If the MAC address of these devices are known, and their actual location will be fixed for an extended period of time, the test point facility can be used to perform a running comparison of their estimated versus actual location, and the results reported back to the user.

When the actual location of a device is specified by the user of the test point facility, and the device’s MAC address is added as an active test point, WCS works in concert with the location appliance’s debugging facility to compare the device’s actual location to its estimated location after each SNMP poll period for that device category. WCS performs these comparisons and calculates the percentage of occurrences where the estimated location of each tracked device is within the baseline performance specification of the location appliance. This information is reported to the WCS user whenever the user highlights the device MAC address in the test point’s MAC address drop-down menu and clicks on the “Analyze” option. When the user has completed tracking a specific MAC address, usage of that device

as a test point can be terminated by highlighting the device MAC address in the test point MAC address drop-down menu and clicking on the “Stop” option. When all tracking has been completed and the accumulated tracked data is no longer needed, all test points and their test point tracking data can be removed by using the “Clear Logs” option.

Additional information regarding the configuration of the test point facility for use can be found at the following URL:

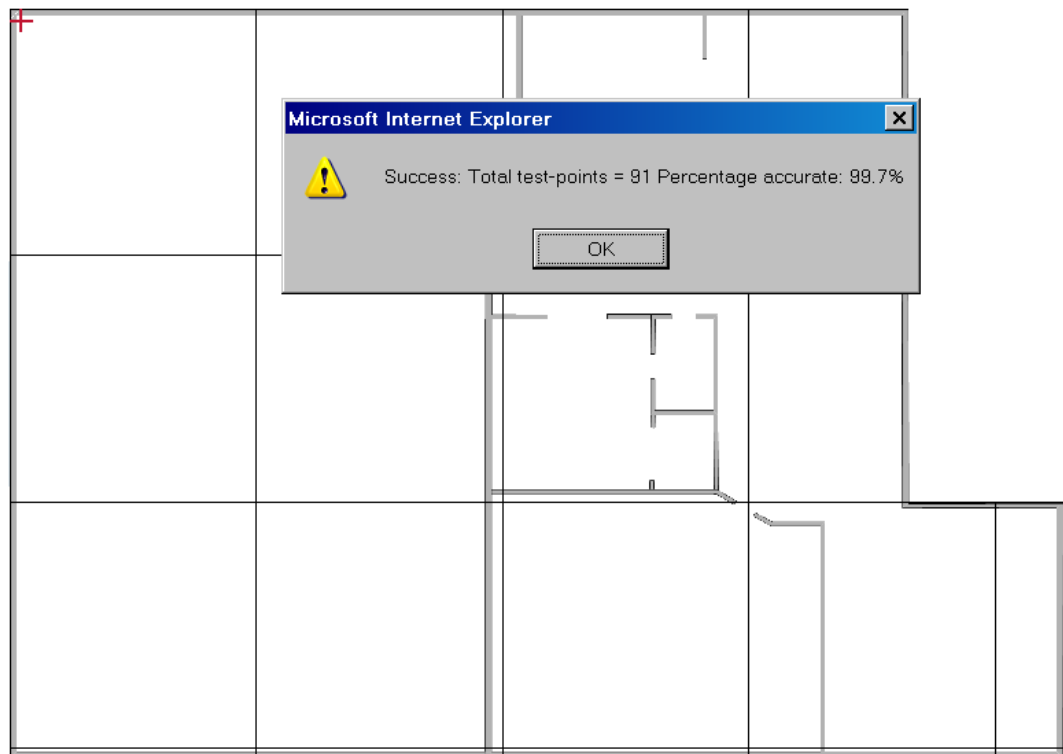
http://www.cisco.com/en/US/docs/wireless/location/2700/3.0/configuration/guide/lacg_ch7.html#wp1066057

Note that before the “Analyze” feature can provide any useful feedback, at least one device category poll period must occur. This is so as to allow the location appliance sufficient time to obtain location information regarding the test points (note that historical location data is not used by the test point facility). It is recommended to allow at least several poll periods to transpire in order to provide the best possible indication of long term location accuracy. Figure 5-48 illustrates an example of the output provided by the “Analyze” feature of the test point facility.

Figure 5-48 Test Point Facility

Position Test Point on Floor 'Main Floor'

Debug In	Location Server	MAC Addr	Horiz	Vert	
Location Server	AeS_Loc1	*** Last Used *** 00:0c:cc:5e:7b:69	0	0	Save
		*** Client List *** 00:40:96:a1:9d:47			Preview
		***** Tag List *****			Analyze
					Stop



2233569

When using the test point facility, keep in mind the following:

- Advanced Debug must be enabled on WCS prior to attempting to use the test point facility. If Advanced Debug is not enabled on WCS, the “Position Test Point” option will not appear in the dropdown menu located at **Monitor > Maps > floor map name**.
- Although it is a recommended best practice to enable the Advanced Debug on the location appliance prior to using the test point facility, it will be enabled automatically whenever a test point is added using the “Save” option.
- The “Preview” option is used to position the red cross hair on the Position Test Point floor map only when the user is directly specifying horizontal and vertical (x,y) coordinates. This is done instead of manually positioning the red crosshair to the test point’s actual location. The preview capability is provided so that graphical indication of the location corresponding to the (x,y) coordinates just entered can be presented as visual confirmation to the user. If the user manually positions the crosshairs to the test point’s location and then clicks on preview, they will not receive any feedback. The preview feature is not intended to display the estimated location of the test point device.
- In [Figure 5-48](#), note that the “Total Test Points” quantity reflected in the output of the analyze command does not indicate the total number of test point devices currently in use. Rather, it indicates the total number of test point location samples taken for the highlighted device MAC address.
- Although multiple devices can be selected on the Position Test Point WCS screen, test points can only be added (using “Save”) one device at a time. The MAC address of the last device used for a save or stop operation is retained and shown under the “Last Used” heading.
- Any test points that have been added using “Save” will no longer be accessible to the user if they log out of WCS and then log back in at a later time. In order to ensure that all such residual test points are cleared prior to beginning a new test session, it is a good idea to issue “Clear Logs” before beginning the definition of a new set of test points.
- Test points that have been added will not survive a reboot or restart of the appliance. Collected test data will still be resident on the location appliance, however, it will not be possible to add further test point data to that already collected. In order to proceed with further test point data collection and analysis, it is recommended that the device MAC address be stopped or the “Clear Logs” option be used.
- The “Clear Logs” option clears *all* test point data logs and resets any assigned test point devices. Keep in mind that clearing the logs will delete not only the test points created by the current user, but those test points created by *any* other users that are logged in and authorized to make use of the test point facility as well.

In software Release 4.1, the “Analyze” option performs its calculations in regard to a single device MAC address at a time. If multiple device MAC addresses are added as test points, examine the test point data for each device MAC address individually by highlighting the MAC address of the device in the drop-down selector and then clicking on “Analyze”.



CHAPTER 6

RFID Tag Considerations

This chapter has the following main sections:

- [RFID Tag Technology, page 6-1](#)
- [Using Wi-Fi RFID Tags with the Cisco UWN, page 6-15](#)
- [Tag Telemetry and Notification Considerations, page 6-27](#)
- [Chokepoint Considerations, page 6-31](#)

RFID Tag Technology

The majority of RFID tags produced today are *passive* RFID tags, comprised basically of a micro-circuit and an antenna. They are referred to as passive tags because the only time at which they are actively communicating is when they are within relatively close proximity of a passive RFID tag reader or *interrogator*.

Another type of common RFID tag in the marketplace today 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 tables shown in [Figure 6-1](#) provide a quick reference of common comparisons between active and passive RFID tags. Within these basic categories of RFID tags can be found subcategories such as *semi-passive* RFID tags.



Note

The terms *beacon* and *beaconing* have been used in the RFID industry for some time, predating the establishment of the formal 802.11 standards. When an active RFID tag periodically beacons, it is simply transmitting a tag message (much like any other messages the tag might send) at a set interval. Despite the use of similar terminology, this should *not* be confused with an 802.11 Beacon. An 802.11 Beacon is a management frame that the 802.11 access point (or the beacon sender in an IBSS) transmits to provide time synchronization and PHY-specific parameters in order to facilitate mobile stations locating and identifying a BSS or IBSS.

Figure 6-1 Active and Passive RFID Comparison

	Active RFID	Passive RFID
Tag Power Source	Internal to tag	Energy transferred from the reader via RF
Tag Battery	Yes	No
Availability of Tag Power	Continuous	Only within field of reader
Required Signal Strength from Reader to Tag	Very Low	Very High (must power the tag)
Available Signal Strength from Tag to Reader	High	Very Low

	Active RFID	Passive RFID
Communication Range	Long range (100m or more)	Short or very short range (3m or less)
Sensor Capability	Ability to continuously monitor and record sensor input; data/time stamp for sensor events	Ability to read and transfer sensor values only when tag is powered by reader; no date/time stamp
Data Storage	Large read/write data storage (128KB) with sophisticated data search and access capabilities available	Small read/write data storage (e.g. 128 bytes)

190588


Recent market developments have brought yet another category of RFID tag into the spotlight. Known as hybrid or *multimode* tags, these combine several different tag technologies into a versatile package that can be tracked by one or more location technologies. Multimode RFID tags are typically low power, small form factor devices that allow a single physical tag to assume multiple personalities and perform tasks that previously would have required several individual physical tags to be attached to the asset. A multimode tag, for example, may combine multiple active tag subcategories along with a passive tag into a single homogenous product.

Passive RFID Tags

Passive RFID tags typically do not possess an onboard source of power. Instead, the passive RFID tag receives its power from the energizing electromagnetic field of an RFID reader (or interrogator). The energy coupled from the electromagnetic field undergoes rectification and voltage multiplication in order to allow it to be used to power the passive tag's microelectronics. In the typical passive RFID tag design, the tag cannot communicate with host applications unless it is within the range of an RFID reader.

Interrogators come in many forms, with two common examples being handheld reader-interrogators (shown on the left in Figure 6-2) and large stationary models capable of reading many tags simultaneously as they pass (shown in the center of Figure 6-2). Embedded sub-miniature passive RFID readers and tags (shown on the right in Figure 6-2) can be used in applications requiring immediate action verification. Examples of this might include immediate verification of proper supply-line hose connections. In these types of applications, passive RFID tags and microreaders embedded into hose plugs and receptacles ensure that the proper supply hoses are connected to the proper material sources at all times. Should an incorrect connection be made, the mismatch is detected and the system refuses to open an electromagnetic flow control.


Figure 6-2 *Passive RFID Interrogators*



190589

Passive RFID tags (shown in [Figure 6-3](#)) consist of a coil and a microcircuit that includes basic modulation circuitry, an antenna, and non-volatile memory.

Figure 6-3 *Passive RFID Tags*



190590


Passive RFID tags vary in how they communicate data to RFID readers and how they receive power from the RFID reader's inductive or electromagnetic field. This is commonly performed via two basic methods:

- Load modulation and inductive coupling in the near field—In this approach (see [Figure 6-4](#)), the RFID reader provides a short-range alternating current magnetic field that the passive RFID tag uses for both power and as a communication medium. Via a technique known as *inductive (or near-field) coupling*¹, this magnetic field induces a voltage in the antenna coil of the RFID tag, which in turn powers the tag. The tag transmits its information to the RFID reader by taking advantage of the fact that each time the tag draws energy from the RFID reader's magnetic field, the RFID reader itself can detect a corresponding voltage drop across its antenna leads. Capitalizing on this phenomenon, the tag can communicate binary information to the reader by switching ON and OFF a load resistor to perform *load modulation*. When the tag performs load modulation, the RFID reader detects this action as amplitude modulation of the signal voltage at the reader's antenna. Load modulation and inductive coupling can be found among passive RFID tags using frequencies from 125 to 135 kHz and 13.56 MHz. Limitations that exist with regard to the use of such low frequencies include the

1. A technique based on Faraday's principle of magnetic induction.


necessity to use larger antennas, low data rate and bandwidth and a rather dramatic decay in the strength of the electromagnetic field ($1/r^6$), where r represents the distance between a low frequency interrogator and a passive RFID tag.

Figure 6-4 *Passive Tag Load Modulation*



- Backscatter modulation and electromagnetic coupling in the far field—In this approach (shown in [Figure 6-5](#)), the RFID reader provides a medium-range electromagnetic field that the passive RFID tag uses for both power and a communication medium. Via a technique known as *electromagnetic (or far-field) coupling*, the passive RFID tag draws energy from the electromagnetic field of the RFID reader. However, the energy contained in the incoming electromagnetic field is partially reflected back to the RFID reader by the passive tag antenna. The precise characteristics of this reflection depend on the load (resistance) connected to the antenna. The tag varies the size of the load that is placed in parallel with the antenna in order to apply amplitude modulation to the reflected electromagnetic waves, thereby enabling it to communicate information payloads back to the RFID reader via *backscatter modulation*. Tags using backscatter modulation and electromagnetic coupling typically provide longer range than inductively coupled tags, and can be found most commonly among passive RFID tags operating at 868 MHz and higher frequencies. Far field coupled tags typically provide significantly longer range than inductively coupled tags, principally due to the much slower rate of attenuation ($1/r^2$) associated with the electromagnetic far-field. Antennas used for tag employing far field coupling are typically smaller than their inductively coupled counterparts.

Figure 6-5 *Passive Tag Backscatter Modulation*



Note that neither of these two techniques allows passive RFID tags to communicate *directly* with 802.11 infrastructure access points. All communication from the passive RFID tag occurs via the RFID reader.

Passive RFID tags are less costly to manufacture than active RFID tags and require almost zero maintenance. These traits of long-life and low-cost make passive RFID tags attractive to retailers and manufacturers for unit, case, and pallet-level tagging in *open-loop* supply chains. Open-loop supply chains typically allow little to no regulation of whether RFID tags leave the control of the tag owner or originator. Because of their dependence on external reader energy fields and their low reflected power output, passive RFID tags have a much shorter read range (from a few inches for tags using load modulation up to a few meters for those using backscatter modulation) as well as lower read reliability when compared to active RFID tags.

The passive RFID tag is available commercially packaged in a wide variety of designs, from mounting on a simple substrate to creating a classic “hard” tag sandwiched between adhesive and paper (commonly referred to as an RFID “smart” label). The form factor used depends primarily on the application intended for the passive RFID tag and can represent the bulk of the passive RFID tag cost.


Semi-Passive RFID Tags

Semi-passive RFID tags overcome two key disadvantages of pure passive RFID tag designs:

- The lack of a continuous source of power for onboard telemetry and sensor asset monitoring circuits.
- Short range.

Semi-passive tags differ from passive tags in that they use an onboard battery to provide power to communication and ancillary support circuits, such as temperature and shock monitoring. It is interesting to note that although they employ an onboard power source, semi-passive RFID tags do not use it to directly generate RF electromagnetic energy. Rather, these tags typically make use of backscatter modulation and reflect electromagnetic energy from the RFID reader to generate a tag response similar to that of standard passive tags (see [Figure 6-6](#)). The onboard battery is used only to provide power for telemetry and backscatter enabling circuits on the tag, not to generate RF energy directly.

Figure 6-6 Backscatter Modulation in Semi-Passive RFID Tags



Semi-passive RFID tags operating in the ISM band (shown in [Figure 6-7](#)) can have a range of up to 30 meters with onboard lithium cell batteries lasting several years. Range is vastly improved over conventional passive RFID tags primarily because of the use of a backscatter-optimized antenna in the semi-passive design. Unlike a conventional backscatter-modulated passive RFID tag, the antenna contained in a semi-passive tag is dedicated to backscatter modulation and there is no dependence on the semi-passive RFID tag antenna to be a reliable conduit of power for the tag. Therefore, the semi-passive tag antenna can be optimized to make most efficient use of the backscatter technique and provide far better performance than purely passive RFID tag antenna designs.

Figure 6-7 **Semi-Passive RFID Tags**



Several varieties of semi-passive RFID tags exist, with and without onboard NVRAM, real time clocks, and various types of environmental sensors. Semi-passive RFID tags also support interfaces to tamper indicators, shock sensors, and so on. Common applications of semi-passive RFID tags include but are not limited to vehicle asset tracking, security access systems, supply chain automation, cold storage management, and hierarchical asset tracking systems.

Active RFID Tags

Active tags are typically used in real-time tracking of high-value assets in *closed-loop* systems (that is, systems in which the tags are not intended to physically leave the control premises of the tag owner or originator). Higher value assets can usually justify the higher cost of the active tag, and presents strong motivation for tag reuse. Medical equipment, electronic test gear, computer equipment, reusable shipping containers, and assembly line material-in-process are all excellent examples of applications for active tag technology. Active RFID tags (see [Figure 6-8](#)) can provide tracking in terms of *presence* (positive or negative indication of whether an asset is present in a particular area) or real-time location. Active RFID tags are usually physically larger than passive RFID tags. Most RTLS systems are based on the use of active RFID tag technology.

Figure 6-8 Active RFID Tags



Active tags can contain 512 KB or more of RAM, which enables the active tag to store information from attached assets for transmission at the next beacon interval or when polled. This large memory capacity also makes active RFID preferable to passive RFID in situations when the RFID tag cannot simply be used as a “license plate” or reference, to enable an immediate lookup in a host database. A good example of this might be a remote military installation where a host database may or may not be available at all times. By storing critical asset data directly on the tag itself, this information can be retrieved directly from the tag and used regardless of the availability of the host system.

Active RFID tags can be found operating at frequencies including 303, 315, 418, 433, 868, 915, and 2400 MHz with read ranges of 60 to 300 feet. Active RFID tag technology typically display very high read rates and read reliability because of their higher transmitter output, optimized antenna, and reliable source of onboard power. Active RFID tag cost can vary significantly depending on the amount of memory, the battery life required, and whether the tag includes added value features such as onboard temperature sensors, motion detection, or telemetry interfaces. The durability of the tag housing also affects price, with the more durable or specialized housings required for specific tag applications coming at increased cost. As with most electronic components of this nature, prices for active tags can be expected to decline as technological advances, production efficiencies, and product commoditization all exert a downward influence on market pricing.

Beaconing Active RFID Tags

Beaconing active RFID tags are used in many RTLS systems and are primarily useful when the location of an asset needs to be tracked anywhere and anytime via the use of location receivers. With a beaconing active RFID tag, a short message payload containing the unique identifier of the RFID tag is emitted at pre-programmed intervals. This interval is programmed into the tag by the tag owner or user, and it can be set appropriately depending on how often tag RSSI updates are required. A shorter tag transmission interval typically results in shorter tag battery life but may improve tag location accuracy in some cases, since tag RSSI is reported more often. Longer tag transmission intervals increase tag battery life but as tag RSSI is reported less often, the frequency of location update will be less.

802.11 Active RFID Tags

802.11 (Wi-Fi) active RFID tags (shown in Figure 6-9) are designed to operate in the unlicensed ISM bands of 2.4 to 2.4835 GHz or 5.8 to 5.825 GHz. Currently manufactured 802.11 Wi-Fi active RFID tags available at publication are limited to 2.4 GHz.

These tags exhibit the characteristics of active RFID tags, but also comply with applicable IEEE 802.11 standards and protocols. Wi-Fi RFID tags can readily communicate directly with standard Wi-Fi infrastructure without any special hardware or firmware modifications and can co-exist alongside Wi-Fi clients such as laptops, VoWLAN phones, and so on. When powered on, assets equipped with 802.11

Wi-Fi client radios can be tracked natively without the need to have an asset tag attached. Other assets lacking an internal 802.11 Wi-Fi client radio can be tracked via a physically attached 802.11 active RFID tag. A physically attached 802.11 active RFID tag also makes it possible to use the location-aware Cisco UWN to track assets with integrated Wi-Fi client radios when those radios are powered off.

Figure 6-9 802.11 Wi-Fi Active RFID Tags




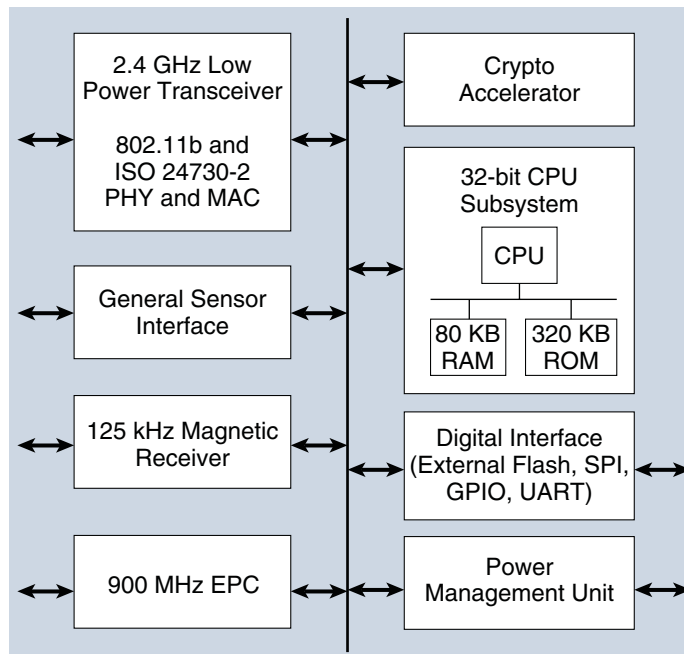
Multimode RFID Tags

As mentioned previously, transponder active RFID tags offer the combination of a primary tag operational mode with a secondary method of communication that can be used for a plethora of added value functions, such as activation, deactivation, behavior modification and so on. This type of tag has been used for quite some time in highway toll plaza applications, for example, where tags are triggered to transmit when in proximity of high speed activators, thereby triggering a debit to the user's account for the toll charge.

A relatively new development has been the introduction of *multimode* RFID tags that leverage multiple location technologies. Multimode tags offer the functional equivalent of having assets equipped with several individual tags in one physical package. This can be very useful when assets must travel outside of a single enterprise closed loop system into other systems, where the same type of location tracking technology may not be in use. For example, consider the case where reusable shipping containers must be tracked at a manufacturer, a distributor and a retailer using a combination of ISO24730-2 TDoA, 802.11 Wi-Fi Active RFID and passive RFID. A multimode tag could offer all three of these technologies in a single small form factor, low power draw package. Such a device may also include the capability to use tag magnetic signaling proximity communication devices as well. This can offer distinct advantages in terms of management, maintenance and overall ease of deployment, especially when compared to equipping assets with three or more physically separate RFID tags.

Multimode tags of this nature have been made much more feasible by the availability of highly integrated tag OEM silicon that combines two or more distinct RFID tag technologies into a single chip or chipset. This is exemplified by the G2C501 from G2 Microsystems (shown in Figure 6-10), which is a complete Wi-Fi system-on-chip (SoC) that includes 802.11b Wi-Fi active RFID, 900 MHz EPC Global Gen 1 Class 0 passive RFID, 2.4 GHz ISO24730-2 TDoA, a 32-bit CPU, crypto accelerator, real-time clock and sensor interfaces.

Figure 6-10 G2C501 RFID System-On-A-Chip (SoC)



The use of highly integrated tag silicon offers many advantages to the tag vendor, including:

- Small form factor
- Low power consumption
- Well documented software and hardware interfaces
- Flexible support for multiple location technologies

A good example of a multimode tag that capitalizes on such capabilities is the WhereNet IV asset tag from WhereNet Corporation (<http://www.wherenet.com>), shown in the lower left hand quadrant of Figure 6-9. The WhereNet IV combines a Cisco Compatible Extensions compliant 802.11 Wi-Fi active tag implementation along with 125kHz magnetic signaling and ISO 24730-2 capabilities in a small, highly integrated design.

Chokepoint Triggers

Chokepoint triggers are proximity communication devices that trigger asset tags to alter their configuration or behavior when the asset tag enters the chokepoint trigger's area of operation. This alteration could be as simple as causing the asset tag to transmit its unique identifier, or more complex, including causing the tag to change its internal configuration or status. One of the prime functions of a chokepoint trigger is to stimulate the asset tag such that it provides indication to the RTLS that the tag


has entered (or exited) the confines of an area known as a *chokepoint*. Chokepoints are tightly defined physical areas (such as entrances, exits or other types of constrictions) that provide passage between connected regions. Figure 6-11 illustrates some common examples of chokepoints.



Note

While *chokepoint triggers* are typically deployed within *chokepoints*, it is often commonplace to hear the term *chokepoint* used to refer to a chokepoint trigger.

Figure 6-11 Common Chokepoint Areas



Outdoor chokepoint locations may include a fenced gate, bridge, toll plaza, or similar passageway. Indoor chokepoint locations includes connecting entrances or exits between:

- A building's interior rooms or floors such as doorways, ramps, gates, stairwells, elevator entrances, and so on).
- Adjacent structures (such as passageways or tunnels) or the interior and exterior of structures (main and auxiliary entrances, loading docks, fire exits, and so on).

Chokepoint triggers can initiate behavioral changes in tags that can immediately alert the location system that the tagged asset has entered or exited the chokepoint area. Due to the comparatively modest range of chokepoint triggers in relation to the overall area covered by an RTLS, the RTLS is able to deterministically localize the asset to the confines of the chokepoint area relatively quickly and with excellent reliability. In addition to displaying the chokepoint area on floor maps, the RTLS can use the

detection of assets within chokepoints to trigger events in external systems. These can include database updates, notification alerts, or alarms. When properly augmented by appropriate application software, chokepoint applications may include:

- Tracking of high value assets—Chokepoint location tracking can help ensure that valuable assets intended for a particular area stay within such areas. If these assets are detected as being removed via entrances or exits, for example, the RTLS is alerted.
- Manufacturing process control—Equipment, parts, and finished products can be precisely tracked as they move between the various production stations. This helps ensure not only that all required process stations are visited, but that they are visited in the proper sequence.
- Inventory control—By strategically equipping all distribution center entrances and exits with chokepoint location tracking capabilities, inventory databases can be automatically updated as product enters or leaves the distribution center.
- Security—The movement of tagged assets can be tracked and monitored to protect against unauthorized removal from the premises or unauthorized movement within the facility itself structure.

Low power, short range chokepoint triggers make it possible to expand usage beyond traditional entry and exit passages. Low output power enables customization of the chokepoint trigger's effective range to better correspond to very small, tightly defined areas such as shelves, racks, storage bins, workstations and patient beds. The movement of assets into or away from such limited areas can be then be precisely monitored (such as the placement or removal of equipment in a rack, for example) in a similar fashion to that of the higher power chokepoint triggers described earlier.

The specific changes in tag behavior that can be enacted by a chokepoint are vendor dependent. Tag behavior modification may include, but are not limited to:

- Immediate tag multicast message transmission
- Tag reactivation
- Tag deactivation
- Tag transmission interval change
- Indicator lamp activation
- Storage of floor or cell identifiers
- Appending of additional messages to tag multicast messages, such as:
 - Chokepoint identification
 - Pre-configured message data
 - Telemetry data

Not every active tag vendor supports the use of chokepoint triggers with their tags. Of those that do, the use of chokepoint triggers tends to be tag vendor specific. Each vendor offering asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification usually supplies chokepoint triggers that are designed specifically for compatibility with those tags. At the current time, chokepoint triggers are not interoperable between asset tags from different manufacturers.

Range may vary between models and manufacturers, with those chokepoint triggers used with asset tags compliant with the Cisco Compatible Extensions for Wi-Fi tags specification typically possessing effective ranges between 10 inches and approximately 25 feet. These products operate using low frequency magnetic signaling. Range tends to be predictable, with excellent penetration of typical building materials and their contents.

Figure 6-12 depicts low frequency, magnetic signaling-based chokepoint trigger devices from AeroScout and WhereNet. AeroScout refers to their chokepoint triggers as *Exciters* and WhereNet refers to their products as *WherePorts*. The AeroScout EX-2000 Exciter and the WhereNet WherePort products are larger footprint models, capable of providing the maximum possible range for large chokepoint areas or room-based presence detection applications. These products are intended for vehicular doorways, gates and other large chokepoint areas, with adjustable ranges that can exceed 20 feet. The compact AeroScout EX-3100 and EX-3200 Exciters are intended for short range use in smaller chokepoints such as doorways, shelves and racks. The range of these products spans from 8 inches to a maximum of 6.5 and 9.75 feet, respectively.

Figure 6-12 AeroScout Exciters and WhereNet WherePorts



Additional information on these products can be found at the following vendor web sites:

<http://www.aeroscout.com/content.asp?page=exciter>

http://www.wherenet.com/products_whereport.shtml



Note

The Cisco WCS is used to define chokepoint triggers to the location-aware Cisco UWN, but cannot be used to configure the chokepoint triggers themselves at this time. This must be accomplished using software provided by the vendor of the chokepoint trigger (the AeroScout Network Exciter Manager (ANEM) and the WhereNet SystemBuilder / WhereWand are two examples). Chokepoint triggers that have been added to WCS without proper configuration by the vendor's chokepoint management software may not function properly.

Once configured, chokepoint triggers can operate in one of two modes:

- An *online mode*, where their status is monitored by software supplied by the chokepoint trigger vendor via an Ethernet or serial data connection.
- An *offline mode*, where the configured chokepoint trigger operates with only a power connection required.

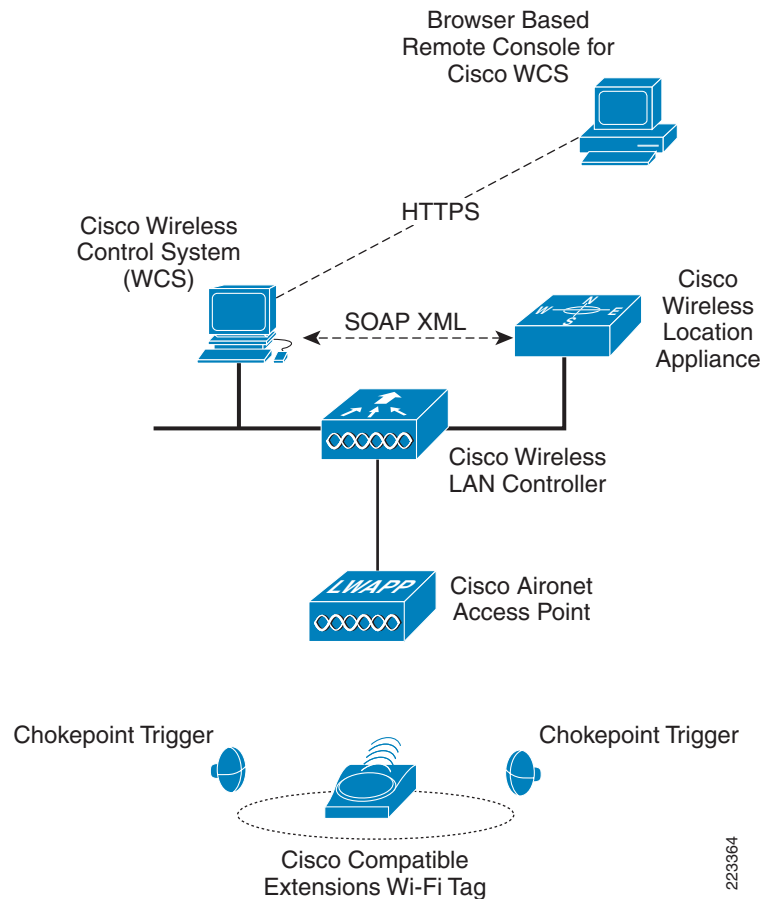
Chokepoint triggers are identified by unique addresses that enables tags receiving their transmission to clearly identify the chokepoint trigger responsible for stimulating them. This identifier is typically the MAC address of the chokepoint trigger for Ethernet-based models, but could be any locally administered and assigned identifier (such as a “Transmit ID” of a WhereNet WherePort). In Release 4.1 of the location-aware Cisco UWN (shown in Figure 6-13), when an asset tag compatible with the Cisco

Compatible Extensions for Wi-Fi Tags specification enters the effective range of a chokepoint trigger, the tag is stimulated by the chokepoint trigger and identifies the source of such stimulation to the location-aware Cisco UWN using a tag multicast frame that is sent via using 802.11. All access points detecting this tag multicast frame forwards it to their registered controller, which in turn results in the generation of LOCP Measurement Notification frames destined for the location appliance.

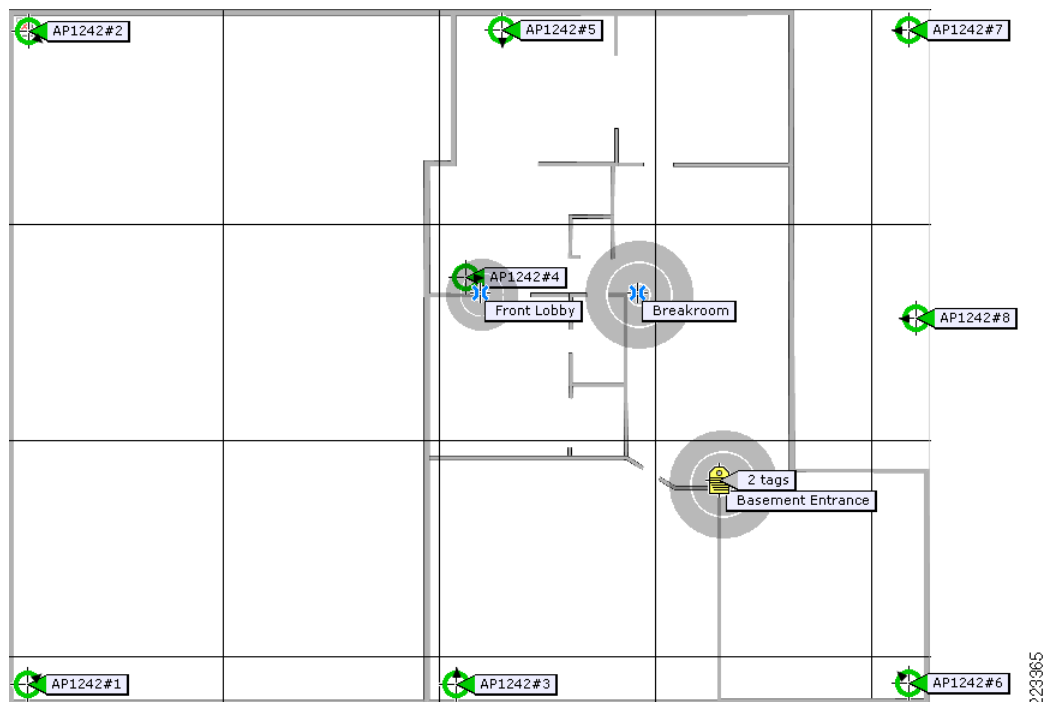
**Note**

Communication between chokepoint triggers and asset tags is unidirectional, from the chokepoint trigger to the asset tag. In software Release 4.1, there is no direct communication between chokepoint triggers and the location-aware Cisco UWN.

Figure 6-13 Location-Aware Cisco UWN with Chokepoint Triggers



The location appliance uses the information provided to it by the LOCP Measurement Notification to indicate that the tag's current location is within the configured range of the specified chokepoint. This information is placed in the appropriate location appliance databases and made available to location clients via the location appliance API. Location clients may display chokepoint location information on floor maps. An example is the WCS floor map shown in [Figure 6-14](#), where we can see two RFID tags located at the chokepoint labeled *Basement Entrance*). The location appliance can also trigger alerts and other asynchronous northbound notifications to WCS and external applications using email, syslog, SOAP, or SNMP traps.

Figure 6-14 WCS Floor Map With Chokepoints

In Release 4.1 of the Cisco UWN software, after a tag has left the range of a chokepoint trigger, the location appliance continues to indicate the tag's location as being within the configured range of the chokepoint trigger until one of the following events occur:

- The tag indicates that it is now out of range of that chokepoint trigger.
- The value configured for the Chokepoint Out of Range Timeout expires (shown in [Figure 6-15](#), default 60 seconds).

After one of these events occur, the location appliance uses RF Fingerprinting to calculate the location of the device until such point that it enters into another chokepoint area and into the stimulation zone of another chokepoint trigger. If the device is then stimulated by a subsequent chokepoint trigger and successfully reports this stimulation to the Cisco UWN, the location appliance then places the tracked device at the location of the new chokepoint.

Figure 6-15 Chokepoint Out of Range Timeout

Location Server > Location Parameters > 'AeS_Loc2'

Location Parameters

Enable calculation time ?	<input type="checkbox"/> Enable
Enable OW Location ?	<input type="checkbox"/> Enable
Relative discard RSSI time ?	<input type="text" value="3"/> minutes.
Absolute discard RSSI time ?	<input type="text" value="60"/> minutes.
RSSI Cutoff ?	<input type="text" value="-75"/> dBm
Smooth Location Positions ?	<input type="text" value="More smoothing (new value weighted less)"/>
Chokepoint Usage ?	<input checked="" type="checkbox"/> Enable
Chokepoint Out of Range Timeout ?	<input type="text" value="60"/> seconds.

223366

Using Wi-Fi RFID Tags with the Cisco UWN

Compatible RFID Tags

An often asked question revolves around whether the Cisco Location Appliance can be leveraged to track RFID tags that already are being deployed by product and durable goods manufacturers as part of a larger business initiative. Often applied *en masse* to manufactured or distributed goods, these tags are most commonly passive RFID designs, but in the case of some durable high-cost goods, active RFID may also be used. In many cases, products and goods are being tagged at the time of production or initial distribution in compliance with mandates set forth by large commercial or governmental entities.

The answer depends on the type of RFID tag being used. As of Cisco UWN software Release 4.1, only 802.11 Wi-Fi active RFID tags (or multimode asset tags containing 802.11 Wi-Fi active RFID capabilities) can communicate directly with Wi-Fi access points (including Cisco Wi-Fi access points). At this time, most commonly available “pure” passive RFID tags or non-Wi-Fi active RFID tags are not capable of communicating with the location-aware Cisco UWN and the Cisco Wireless Location Appliance. Of the available 802.11 Wi-Fi active tag designs currently on the market, not all are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. Non-compliant asset tags from PanGo / InnerWireless and AeroScout Ltd. can be recognized by the location-aware Cisco UWN. However, these tags will not be able to make use of the advanced features in the Cisco Compatible Extensions for Wi-Fi Tags specification and introduced in Release 4.1. Non-compliant asset tags from vendors other than PanGo Networks and AeroScout are not supported for use with the Cisco Wireless Location Appliance.

To determine whether a Wi-Fi active RFID tag is compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification and capable of taking advantage of the advanced features of the location-aware Cisco UWN, the Cisco Compatible Extensions website (http://www.cisco.com/web/partners/pr46/pr147/ccx_wifi_tags.html) should be consulted. A current listing of all tags and tag vendors compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification may be found there.

The listing of the tag and tag vendor on the Cisco Compatible Extensions website indicate that the asset tag has passed stringent validation testing as part of the Cisco Compatible Extensions Program for Wi-Fi tags. The Cisco Compatible Extensions program for Wi-Fi tags allows customers with a location-aware Cisco Unified Wireless Network to benefit from the latest innovation and technology advancements offered by Cisco's technology partners. Registered channel partners may view the guidelines for the Cisco Compatible Extensions Program for Wi-Fi Tags at the following URL:

<http://www.cisco.com/web/partners/downloads/partner/WWChannels/download/wifiguide.pdf>.

In some cases, passive or non-802.11 active RFID reader interrogators may be deployed in an environment that is also serviced by a Cisco LWAPP-enabled wireless network, independently of the location tracking capabilities of the Cisco UWN and the location appliance. These reader/interrogators may be using traditional wired Ethernet as their uplink to the network, or they may have an integrated Wi-Fi client radio (such as the case of portable RFID interrogators like those shown in Figure 6-16). Although it is not possible at this time to track the individual passive RFID tags associated with these portable RFID tag readers using the Cisco location appliance, tracking the portable readers themselves is typically feasible because of their use of industry standard 802.11 client radios. As long as these readers act as standard WLAN clients and authenticate/associate to WLAN SSIDs serviced by controllers defined to the location appliance, they are treated just as other WLAN clients and are indicated on floor maps by a blue rectangular icon.

Figure 6-16 Portable RFID Interrogators with Integrated Wi-Fi Uplink



Using 802.11b Tags in an 802.11g Environment

Another common question that often arises is about the potential performance impact of using an 802.11b asset tag in a network that otherwise consists entirely of 802.11g clients and access points. The crux of such discussions is typically centered around whether or not protection mechanisms (such as RTS-CTS or CTS-to-self) are initiated by the 802.11g network to assure compatibility between the 802.11b asset tags and the 802.11g network.



Note

For an explanation of 802.11g performance, capacity, and protection mechanisms, see the whitepaper entitled *Capacity, Coverage and Deployment Considerations for IEEE 802.11g* at the following URL: http://www.cisco.com/en/US/products/hw/wireless/ps430/products_white_paper09186a00801d61a3.shtml.

A popular point of discussion often revolves around whether these protection mechanisms are initiated upon the introduction of one or more of the following to the all-802.11g wireless infrastructure:

- An 802.11b asset tag that is transmitting tag layer two multicast messages.
- An 802.11b asset tag (acting as a WLAN client) that is issuing probe requests.
- An 802.11b asset tag (acting as a WLAN client) that actively associates.

First and foremost, it should be clearly understood that 802.11b asset tags that transmit tag messages using Layer 2 multicasts (and do not attempt to associate to any WLANs) will *not* cause the initiation of any 802.11g protection modes under any circumstances. This includes asset tags operating in strict compliance with version 1 of the Cisco Compatible Extensions for Wi-Fi tags specification.

Laboratory research and analysis have shown that protection mechanisms are not initiated throughout an entire network of access points if an 802.11b asset tag or WLAN client is simply powered on. In fact, the following are observed:

- A probe request from an 802.11b asset tag that is *not associated* to any access point on a particular channel does not in and of itself cause the initiation of protection mode by an 802.11g access point that detects it.
- Protection mode is not initiated until the 802.11b asset tag successfully associates to either the cell in question or an adjacent cell on the same channel. At that point, the target cell as well as any other cells on the same channel and RF-adjacent to the target cell initiate protection mode.
- Access points that are not on the same channel as the 802.11b asset tag or not RF-adjacent to it does not initiate protection mode.

Some 802.11b asset tags may, as an optional feature, periodically probe and attempt to briefly associate to the wireless infrastructure in order to conduct over-the-air firmware or configuration updates. The observations stated above would apply to these tags, but only during the brief periods during which these extended modes of communication are in use.

Enabling Asset Tag Tracking



Note

Beginning with the Cisco UWN Release 4.1, it is no longer necessary to enable asset tag tracking in WLAN controllers using the **config status rfid enable** CLI command. RFID tag data collection in controllers containing Release 4.1 is now enabled by default.

Enable Asset Tag RF Data Timeout

The *RFID Data Timeout* parameter sets a static time value (in seconds) that must elapse without any access points on the controller detecting an asset tag, before that asset tag is removed from the internal tables of the controller. For general usage, it is recommended that this parameter be set to a minimum of three times (and a maximum of eight times) the longest tag transmission interval found in the general tag population. This should be inclusive of stationary as well as any “in-motion” transmission intervals. The valid range of values for this parameter is 60-7200 seconds and the default value is 1200 seconds.

For example, for a tag with a constant transmission interval of 60 seconds, you may choose to set the RFID data timeout to 480:

```
(Cisco Controller) >config rfid timeout 480
(cisco Controller) >

(Cisco Controller) >show rfid config
RFID Tag data Collection..... Enabled
```

RFID data timeout..... 480 seconds

To ensure proper collection of updated asset tag RSSI from WLAN controllers, it is recommended that the RFID data timeout always be greater than the asset tag polling interval on the location appliance, which is discussed in the next section.

Enable Asset Tag Polling

To use the location appliance for asset tag tracking, SNMP asset tag polling must be explicitly enabled via the Locate > Location Server > Polling Parameters GUI panel. To enable it, use the checkbox indicated by the red rectangle in Figure 6-17.


Figure 6-17 Enabling RFID Tag Polling

The screenshot shows the Cisco Wireless Control System interface. The breadcrumb navigation is 'Location Server > Polling Parameters >'. The 'Polling Parameters' section includes fields for 'Retry Count' (3) and 'Timeout (secs)' (5). Below this is a table with columns 'Enable', 'Polling', and 'Interval (secs)'. The 'Asset Tags' row has an unchecked checkbox in the 'Enable' column, which is highlighted with a red rectangle. Other rows include 'Client Stations', 'Rogues', and 'Statistics'. 'Save' and 'Cancel' buttons are at the bottom.

Enable	Polling	Interval (secs)
<input checked="" type="checkbox"/>	Client Stations	120
<input checked="" type="checkbox"/>	Rogues	600
<input type="checkbox"/>	Asset Tags	120
<input type="checkbox"/>	Statistics	900

The default polling interval value represents the time period between the start of subsequent polling cycles in which the location appliance polls the controller using SNMP. For example, if a polling cycle requires 30 seconds to complete and the polling interval is 300 seconds, polling cycles start every 330 seconds, as shown in Figure 6-18.

Figure 6-18 Polling Interval



Depending on the degree of asset movement, updated tag RSSI information obtained via shorter polling intervals may be translated into more frequent location updates in some cases. However, depending on the time lag between the asset tag polling interval configured on the location appliance and the average transmission interval configured amongst the general tag population, a risk of reduced asset tag polling efficiency may occur. In extreme cases of deployments with a large number of WLAN controllers, a too

short asset tag polling interval could burden both the location appliance as well as the WLAN controllers with almost constant (and often times unproductive) polling. This wastes resources that could have been put to use more productively, and could negatively impact performance.

In general, for a given population of asset tags with the same transmission interval, the most productive and efficient polling is found to occur when the location appliance's asset tag polling interval is configured to be greater than or equal to the asset tag's transmission interval. For example, in a population of 100 asset tags each with a transmission interval of 60 seconds, if the location appliance's asset tag polling interval is left at the default of 120 seconds (twice the tag transmission interval) it is likely that controllers will receive updated RSSI from all 100 tags at least once (and most likely twice) within the 120 second time interval. Setting the asset tag polling interval to 30 seconds in an attempt to increase the frequency of tag location updates might indeed accomplish this goal for some tags, however, overall polling efficiency is likely to decline.

In a population of asset tags that are configured with mixed transmission intervals, a tradeoff typically is required between the desire to acquire frequently updated RSSI information from tags possessing the shortest transmission intervals versus overall polling efficiency for the general tag population. Shorter asset tag polling intervals can be configured to favor tags that transmit multicast frames more frequently, but depending on the number of WLAN controllers deployed, asset tag polling intervals should not be set so short that the location appliance is spending the bulk of its time constantly polling controllers, which could impact performance in an environment with many controllers present. Remember that the speed at which location updates are displayed on location client screens depends not only on the frequency of updates between controllers and the location appliance, but also upon the frequency with which the location client polls the location appliance for updates.

Recording of asset tag location history is disabled by default. If location trending and the analysis of past asset tag location history is desired, location history recording should be enabled via the Location > History Parameters screen, as shown in Figure 6-19. Enable the **Asset Tags** line item and specify the history archival interval between writes of historical data to the database (default is 720 seconds). Note that the recording of location history is not mandatory to perform asset tag tracking, but is often desirable, as it allows the location appliance to “playback” the history of locations the asset tag has visited.

Figure 6-19 Enabling RFID Tag History

The screenshot shows the Cisco Wireless Control System interface. The breadcrumb navigation is 'Location Server > History Parameters >'. The 'History Parameters' section includes the following settings:

- Archive for: 30 days
- Prune data starting at: 23 Hrs, 50 Mins, and also every 1440 minutes.

Enable	History of	Interval (mins)
<input type="checkbox"/>	Client Stations	360
<input type="checkbox"/>	Rogues	720
<input checked="" type="checkbox"/>	Asset Tags	720

At the bottom of the configuration area are 'Save' and 'Cancel' buttons. The 'Asset Tags' row in the table is highlighted with a red border.

190602

Enable Asset Tag Display

For WCS to display the location of asset tags, asset tag display must be explicitly enabled via Monitor > Maps > Campus > Building > Floor, as shown in Figure 6-20. To enable the display of asset tags, make sure that **802.11 Tags** is selected from the dropdown Layers menu. Refresh or reload the WCS floor map page and yellow tag icons is used on the floor map to denote the current location of any detected asset tags.

Figure 6-20 Enabling Display of Asset Tags on WCS

The screenshot displays the Cisco Wireless Control System (WCS) interface. The top navigation bar includes 'Monitor', 'Reports', 'Configure', 'Location', 'Administration', and 'Help'. The current view is 'Maps > Cisco S3 - Site 5 > BLD 14 > 3rd floor'. On the left, a sidebar shows 'Contributing APs' with three checked items: 'sjc14-31b-ap5', 'sjc14-32b-ap6', and 'sjc14-31b-ap3'. Below this is a 'Refresh Heatmap' button and a 'Load' section with a progress indicator and a 'Load Location Server data as old as: 15 Minutes' dropdown. The main area features a 'Layers' menu with '802.11 Tags' selected. Other layers include 'Access Points', 'AP Heatmaps', 'Clients', 'Rogue APs', 'Rogue Clients', 'Grid', 'coverageAreas', 'Markers', and 'Chokepoints'. A 'Save Settings' link is at the bottom of the menu. The floor plan shows several yellow tag icons. Above the map is an 'RSSI Color Lookup' legend ranging from -35 dBm (red) to -90 dBm (purple). To the right of the legend are 'Zoom' (100 %) and 'Refresh' (5 min) controls, along with a 'Full Screen' link. The Cisco logo is in the top left corner, and the number '228867' is in the bottom right corner.

Configuring Asset Tags

In order to communicate with the location-aware Cisco UWN, asset tags must be properly configured for parameters such as channels, transmission interval, and data formats. In this section, we examine the basic parameter settings necessary for AeroScout tags to be recognized by the UWN and properly localized.



Note

AeroScout asset tags are highlighted in this section only as an example of how to configure asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. Keep in mind that each vendor's asset tags require configuration using vendor-specific tools. Users of AeroScout, InnerWireless (PanGo), WhereNet, G2 or other asset tag vendors offering similar products should always consult their vendor's product documentation for appropriate configuration guidelines.

In comparison to the earlier 2.x versions of AeroScout Tag Manager, version 3.x introduces several new features designed to support AeroScout asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification, including the recently introduced AeroScout T3 asset tags.

This section outlines the steps necessary to configure AeroScout asset tags for basic communication with the location-aware Cisco UWN. It does not attempt to serve as a substitute for the much more comprehensive vendor documentation offered by AeroScout in this regard. The following AeroScout documents should serve as the primary reference materials with regard to configuration of AeroScout asset tags using Tag Manager:

- AeroScout Tag Manager Quick Start
- AeroScout Tag Manager 3.0 User Guide

In order to take advantage of the new capabilities introduced by the Cisco Compatible Extensions for Wi-Fi Tags specification, AeroScout asset tags should contain the following tag firmware levels (see [Figure 6-21](#)):

- AeroScout T2—Firmware Release 4.3x or greater
- AeroScout T3—Firmware Release 6.0x or greater

AeroScout asset tags with firmware releases prior to those listed will still interoperate with software Release 4.1 of the location-aware Cisco UWN. However, tags not meeting these specifications will not take advantage of the capabilities introduced by the Cisco Compatible Extensions for Wi-Fi Tags specification that are present in software Release 4.1.

Figure 6-21 AeroScout T2 and T3 Asset Tags



AeroScout asset tags contain both a 2.4 GHz IEEE 802.11b transceiver as well as a low-frequency, short-range 125 kHz magnetic signaling receiver. 2.4 GHz output power is configurable up to a maximum of +19dBm (81mW). During tag configuration, AeroScout asset tags use their 802.11b interface to reply to commands and data received from a programming device known as a *Tag Activator*, which is an Ethernet addressable, low-frequency 125 kHz magnetic signaling transmitter housed in combination with a 802.11b receiver. Tag Activators are designed to be used in conjunction with Windows-based tag configuration software known as *Tag Manager*.

It is important to note that AeroScout asset tags are only capable of receiving information from Tag Activators via their magnetic signaling 125 kHz receiver. AeroScout asset tags are not equipped with a magnetic signaling transmitters, and Tag Activators are not equipped with magnetic signaling receivers. AeroScout asset tags receive commands and data from Tag Activators via magnetic signaling, and respond back to the Tag Manager application confirming those transmissions using their 802.11b capabilities and the 802.11b receiver in the Tag Activator.

The AeroScout Tag Activator (shown in [Figure 6-22](#)) can be powered via 802.3af Ethernet or an external 5VDC power source. The Tag Activator works in conjunction with AeroScout Tag Manager software to configure, program, activate, or deactivate up to 50 AeroScout asset tags simultaneously at a range of up to approximately three feet. The use of a Tag Activator is completely non-intrusive in relation to the AeroScout tag hardware. There are no cables that interconnect the two, and the use of the Tag Activator eliminates disturbing the environmental seal of the tag casing for configuration modifications. Minimal disruption of tag seals is an advantage if the asset tag is intended for use in harsh or wet environments where tight environmental sealing is required.

Figure 6-22 AeroScout Tag Activator



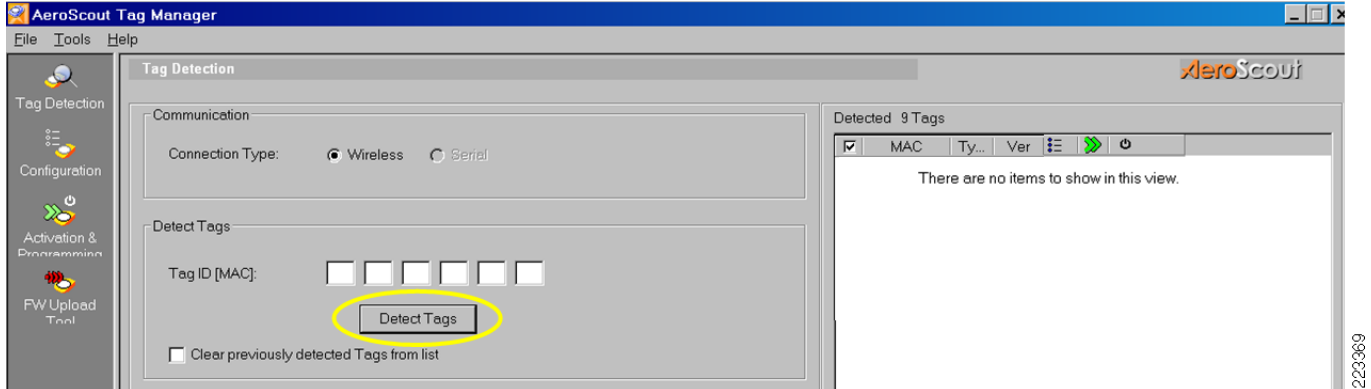
The following AeroScout document should serve as the primary reference with regard to the AeroScout Tag Activator:

- *AeroScout Tag Activator User's Guide*

In order to configure AeroScout T2 or T3 asset tags for basic communication with software Release 4.1, the following steps should be followed:

1. Deploy the AeroScout Tag Activator in accordance with the vendor's recommendations as outlined in the *AeroScout Tag Activator User's Guide*. The AeroScout tag activator may be powered directly from a 802.3af compliant switch or from a non-802.3af switch using the provided AC power supply included with the product. **Spanning tree portfast** should be configured on any Cisco switch port to which the AeroScout Tag Activator is attached to avoid potential instability.
2. Configure the AeroScout Tag Manager to communicate with the Tag Activator as per the vendor's recommendations as outlined in the *AeroScout Tag Activator User's Guide* and the *AeroScout Tag Manager version 3.0, Quick Start Guide*. Ensure that the Tag Activator is properly recognized by the Tag Manager.
3. Place up to 50 AeroScout tags within about three feet of the Tag Activator and detect the tags using the "Detect Tags" feature as shown in [Figure 6-23](#).

Figure 6-23 Detecting Tags using Tag Manager v3.04



- Once the tags have been detected (Figure 6-24), select all tags by clicking on their checkboxes, as shown in the right hand column of the screen depicted in Figure 6-25.

Figure 6-24 Successful Tag Detection using Tag Manager v3.04

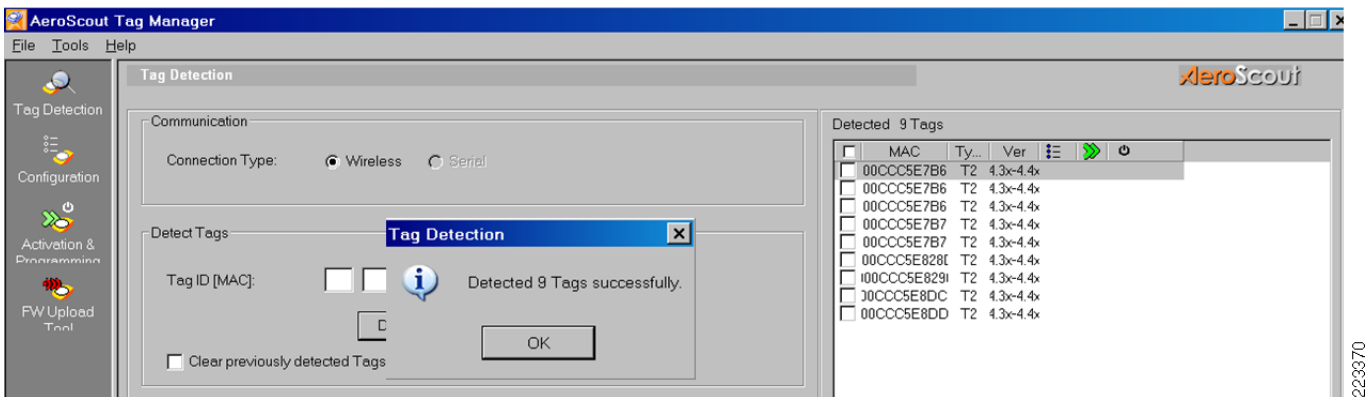
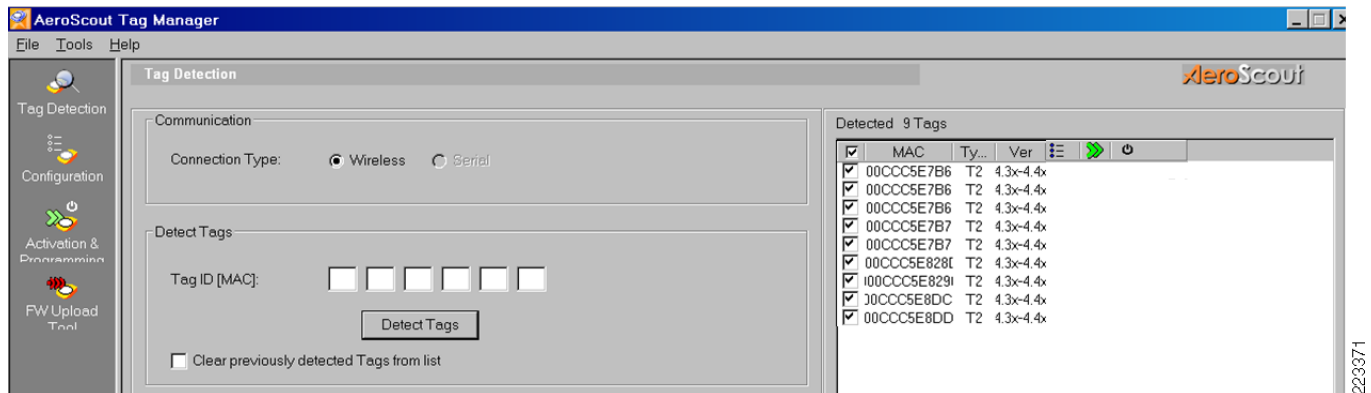


Figure 6-25 Selecting Tags to Configure




- Select the configuration option from the left hand column of the Tag Detection menu, which yields the Tag Configuration menu (shown in Figure 6-26).

**Note**

When making minor modifications to preconfigured tags, it is recommended that the current configuration of the tag be imported into Tag Manager and used as a configuration template, with any modifications then applied to that configuration. The result can then be applied to one or more tags. To do this, after selecting the **Configuration** menu option, place the mouse cursor over the tag that you would like to use as a template. Right click, and select **Get Tag Configuration**, respond **Yes** when asked to proceed.

Figure 6-26 Tag Manager 3.04 Configuration Panel



6. Configure each parameter subcategory for basic operation of T2 or T3 tags with the Cisco UWN software Release 4.1. If you have selected both T2 and T3 tags, note that only the configuration options that apply to *both* tag models are available. Once all parameters in a configuration group have been configured, they may be applied to the selected tags by clicking on the **Apply** button that appears within each group. Alternatively, you may delay applying changes until all groups have been configured (use the **Apply Multiple Configuration** option shown at the bottom of [Figure 6-26](#)). All parameters selected are applied to all selected asset tags and will override any other values that may be present.

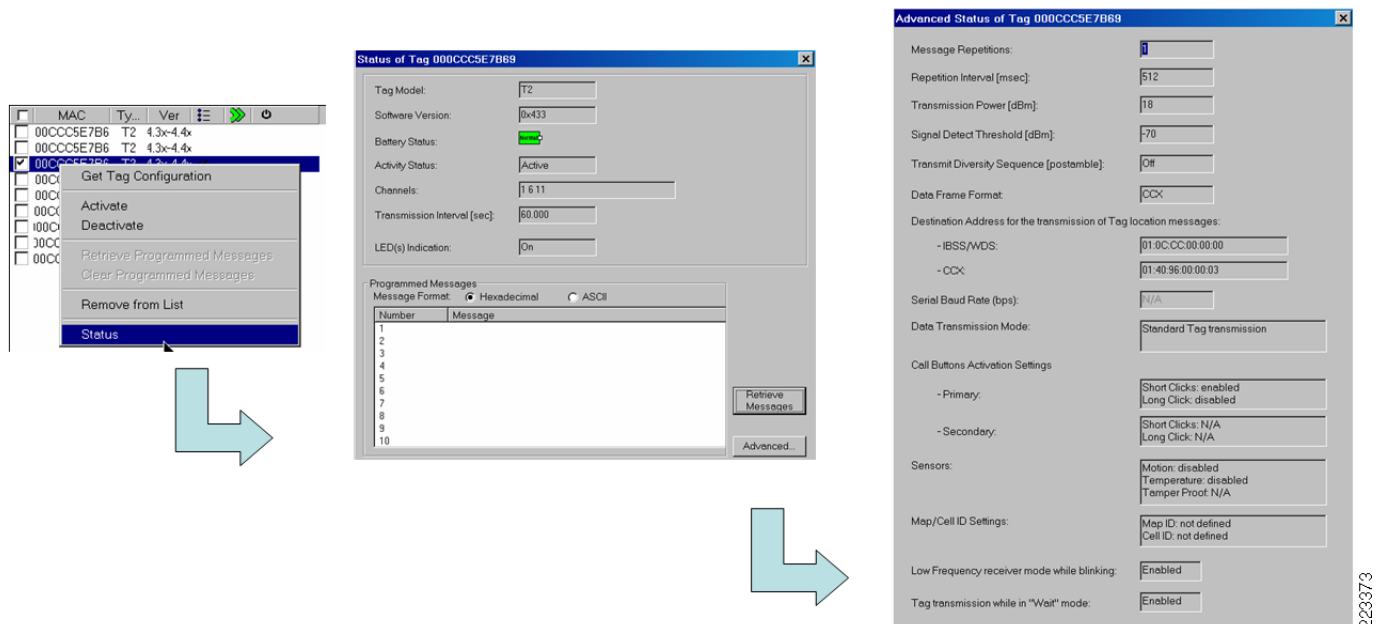
- a. General Parameters:

- Channel Selection—It is recommended that tags be configured for the standard set of 802.11b non-overlapping channels, typically channels 1, 6 and 11 (or otherwise depending on your regulatory domain).

- LED Indication—In most cases, it is useful to have visual indication of when the tag is using its communication interfaces. In cases where there are reasons why such indication is undesirable, such as in a light sensitive, security or other “stealth” application, the LED can be disabled.
 - Transmission Interval When Not In Motion—Select an appropriate tag transmission interval for your asset tagging application, in seconds or milliseconds. Typically tags are configured to transmit less frequently when stationary using this parameter setting as compared to when they are in motion. In-motion transmission intervals are set using the Motion Sensor category settings.
- b. Transmission Parameters:
- Message repetitions—Standard operation for the AeroScout tag is to transmit a single multicast transmission on all defined channels. This parameter controls the number of times each transmitted message is repeated, per channel. It is generally recommended that this parameter be raised from the default value of one to a value of three. Doing this helps protect against lost tag transmissions, which results in lost RSSI readings. Lost RSSI readings is a confirmed cause of degraded location accuracy, especially in environments where there is a significant likelihood of tag transmissions being interfered with or dropped due to congestion or interference. Avoid configuring an excessive number of message repetitions, as there are few conditions where a message repetition factor greater than 3 would be truly required. The setting of three message repetitions works very well for the majority of environments. Setting this parameter above a value of 5 is typically not considered necessary.
 - Message Repetitions Interval—The delay between subsequent message repetitions on the same channel, specified as either 128, 256 or 512 milliseconds. The default value is 512 milliseconds.
 - Transmission Power (dBm)—The default value for transmission power is typically +18dBm on T2 model AeroScout asset tags. The location-aware Cisco UWN is capable of discerning the transmission power used by tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification.
 - Data Rate—Data rates of 2 Mbps can only be specified for T3 tags. Although the message payloads and frame sizes associated with asset tags are very small, the use of a faster transmission speed can allow T3 tags to transmit their payloads faster and free the channel for use by other stations sooner. This can also reduce battery consumption since each frame’s transmission time is shorter.
 - Data Frame Format—This parameter should be changed from the default value of **IBSS** to **CCX**.
 - Destination Address—This value must be specified as 01:40:96:00:00:03 for use with software Release 4.1 and later releases.
- c. Data Transmission Mode Parameters:
- Normal Tag Transmission (without additional message)—Select this parameter unless you have valid reasons to configure it otherwise. For example, the location client you are using in conjunction with your asset tags may be able to process additional stored messages on your tag, sent as part of tag payloads, or you may be using an AeroScout T2 telemetry tag that allows for telemetry to be read directly from sensors onboard custom-integrated host peripheral devices.
- d. Supplementary Settings:
- CCX Options—Transmit Out of Range Chokepoint Group should be enabled.
- e. Call Buttons Primary —Configure these options if you wish to use call button signaling (Panic Button alerting) with software Release 4.1.
- Short Clicks (button depression that last less than 2 seconds):
 - Enable Short Clicks should be checked

- Number of Short Clicks: 1
 - Tag Reaction Parameters: Send Standard Tag Transmission
 - Message Repetition: 1
 - Long Clicks (button depression that lasts at least 2 seconds):
 - Enable Long Clicks should be checked
 - Number of Long Clicks: 1
 - Tag Reaction Parameters: Send Standard Tag Transmission
- f. Call Buttons - Secondary—These are identical options to those listed for “Call Buttons – Primary” but are only available if you are using T3 asset tags.
- g. Sensors:
- Motion—These options can be used to enable the on-board motion sensor if desired.
 - Temperature—These options can be used to enable on-board temperature sensors if desired. Note that the on-board temperature sensor is not supported in T2 tags with v4.3x firmware.
 - Tamper—This option can be enabled for T3 tags only. Enabling this option allows tag tamper indication to be sent to the Cisco UWN.
7. In some cases, the existing configuration of an AeroScout asset tag may be in question and need verification. Using Tag Manager v3.04, this is a straightforward process. Simply right-click on any detected tag and click on **Status** from the pop-up menu. This brings up a listing of basic tag configuration parameters, with further detail available by selecting Advanced Configuration as shown in Figure 6-27.

Figure 6-27 Retrieving The Configuration of a Single Tag



223373

The preceding quick, seven-step configuration guide is just a short synopsis of the required steps to configure and activate AeroScout tags for use with the Cisco UWN software Release 4.1. Refer to the *AeroScout Tag Manager v3.0 User's Guide* for more detailed information as well as information on several other useful configuration options in the Tag Manager.

Tag Telemetry and Notification Considerations

Beginning with software Release 4.1, the location aware Cisco UWN will recognize tag telemetry and high priority notifications transmitted by Wi-Fi Tags specification may transmit tag telemetry and high-priority notifications to the location-aware Cisco UWN. This information is passed from WLAN controllers to the Cisco Wireless Location Appliance using the Location Control Protocol (LOCP), which is described in [Cisco Location Control Protocol \(LOCP\), page 3-36](#).

This section provides initial best practice recommendations and other information and should be kept in mind when designing solutions that are dependent on telemetry and high-priority notification functions found in Cisco UWN software Release 4.1.

Deploying Tag Telemetry

Active RFID tags supplied by tag vendors in compliance with the Cisco Compatible Extensions for Wi-Fi Tags specification may include the ability to accept telemetry data from onboard sensors or from sensors integrated into the asset to which the tag is attached. If configured to do so, these active RFID tags can pass this telemetry data as part of the tag transmissions that are sent to the Cisco UWN at periodic transmission intervals, or when entering into the stimulation zone of chokepoint triggers.

For example, an asset tag connected to the fuel level sensor of a forklift may be able to pass fuel level telemetry via the Cisco UWN to the location appliance and its location clients (which could include WCS and third party location clients). The ability of the asset tag to perform these telemetry functions is dependent upon the asset tag manufacturer, and typically requires the appropriate level of integration and physical connectivity between the tag and sensors found aboard the attached asset. Note that some asset tags are available with their own onboard sensors, which can measure certain ambient environmental characteristics (such as temperature and humidity) external to tagged assets without any dependence on embedded sensors.

Onboard tag sensors, for example, might be appropriate where the primary concern surrounds general environmental conditions effecting both the asset tag as well as the asset to which it is attached. Thus, an asset tag equipped with onboard temperature sensors would be appropriate in detecting whether an attached asset was incorrectly stored in temperatures outside recommended ranges. Embedded sensors within the asset itself would be more appropriate when the goal is to alert the system administrator to an internal condition resulting from improper use that could result in costly damage to the asset if not addressed promptly. A good example of this might be an engine providing indication of an insufficient internal lubrication, which could result in costly repairs.

As described in the section entitled [Asset Tag Telemetry Using LOCP, page 3-38](#), beginning with the Cisco UWN software Release 4.1 all tag telemetry sent by tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification is aggregated by WLAN controllers and passed to the Cisco Wireless Location Appliance. In software Release 4.1, LOCP uses a polled mechanism to collect tag telemetry after the fact, the timing of which is tied to the traditional SNMP polling mechanism used to gather asset tag RSSI information. The location appliance updates the telemetry information for each asset tag in its databases with that received from the most recently responding WLAN controller that has

included telemetry information for that specific tag's MAC address. If archiving of tag historical information has been enabled on the location appliance, tag telemetry information is included along with other tag information (shown in Figure 6-28).

Figure 6-28 Archive Playback of Tag Telemetry and "Emergency" Data

Aeroscout Tag 00:0c:cc:5c:05:13 -- Select a command --

Asset Name	Asset Group
Asset Category	MAC Address 00:0c:cc:5c:05:13

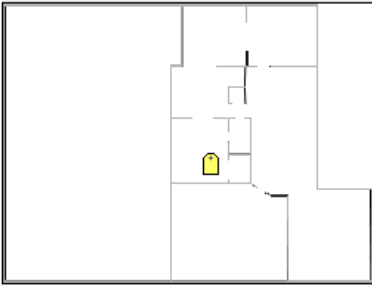
From : Tue Apr 24 23:01:53 EDT 2007
To : Tue May 15 19:13:23 EDT 2007

Time Stamp	Floor	Battery Status
17 Sat Apr 28 10:03:04 EDT 2007	Lab Annex #2 Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %
18 Fri Apr 27 21:09:29 EDT 2007	Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %
19 Fri Apr 27 19:09:29 EDT 2007	Alpharetta Campus>AP1242 Building>Test Lab Annex #2	80 %

Change selection every

Location

Location Calculated	Fri Apr 27 19:08:25 EDT 2007
Floor	Alpharetta Campus>AP1242 Building>Test Lab Annex #2



[Enlarge](#)

Tag Statistics

Data Collected	Fri Apr 27 19:07:37 EDT 2007
Bytes received	343228
Packets received	6601

Telemetry Data

MOTIONPROB : Movement

Emergency Data

Reason:	Tampering
Tamper State:	Active

Tag Properties

Data Collected	Fri Apr 27 19:07:25 EDT 2007
Controller	10.1.96.18
Battery Status	80 %

223374

The default configuration of some active RFID tags may provide for transmitting only one tag transmission per channel per transmission interval. While this setting can help optimize the battery life of the tag in some cases, this single transmission per channel may not always be successfully detected by the expected number of access points, especially in RF-noisy or congested environments. This can result in missing RSSI readings, which can cause location inaccuracy.

Therefore, in such environments it is recommended that tags be configured to transmit multiple transmission repetitions per channel at each transmission interval, which should aid in improving tag detection and location accuracy as well as increasing the reliability of tag telemetry as well. It is recommended that the tag vendor's configuration software should be used to set the number of tag transmissions to three (but not more than five) per channel per transmission interval.

Although it is unlikely that LOCP telemetry collection will burden modern wired and wireless networks, nevertheless it is good practice for the network designer to understand the nature of the traffic that can be expected in their designs. The following traffic and frame size information has been observed during LOCP telemetry testing in support of this document:

- *Echo Request*—Sent periodically by the location appliance to each defined WLAN controller based on the configuration of the Echo Interval parameter (Location Servers > Advanced > LOCP Parameters). LOCP Echo Request Ethernet frames are 100 bytes in length and are transmitted to TCP destination port 16113.
- *Echo Response*—Sent periodically by each WLAN controller in response to an Echo Request (see above). Like Echo Requests, LOCP Echo Response Ethernet frames are 100 bytes in length.
- *Information Request*—Sent periodically by the Location appliance to each WLAN controller to request information. LOCP Information Request Ethernet frames are 106 bytes in length are transmitted to TCP destination port 16113. LOCP Information Requests are the primary mechanism used in software Release 4.1 to conduct LOCP polling.
- *Information Response*—Sent periodically by each WLAN controller in response to the receipt of a LOCP Information Request frame (LOCP Polling). The basic size of a LOCP Information Request Ethernet frame for a controller that has not detected any tags is 113 bytes. If one tag is detected, this frame size will increase to 144 bytes and for two tags it will increase to 175 bytes (these frame sizes do not include any telemetry data). Frame sizes will increase based on the number of tags currently active in the controller's database as well as the amount of telemetry that has been collected. Support for fragmentation and reassembly of combined tag payloads is inherently to LOCP.

To ensure proper LOCP operation between the location appliance and any WLAN controllers defined to it, ensure that port 16113 is not blocked by any firewalls or other security devices.

When designing solutions that will rely on the reporting and collection of tag telemetry with Release 4.1, there are a few considerations that should be kept in mind:

1. *Telemetry Timing*—Since in Release 4.1 telemetry is aggregated on a per-tag basis by WLAN controllers and passed to the location appliance only during a periodic LOCP polling cycle, users of software Release 4.1 should not rely on the receipt of tag telemetry to be real-time in nature. It is reasonable to expect that there will be a delay between the time the tag sends the telemetry information and the time it is updated in the location appliance database and made available to location clients.
2. *Northbound Asynchronous Notifications*—In Release 4.1 of the location-aware Cisco UWN, the location appliance does not issue asynchronous northbound notifications (in the form of email, SNMP, SOAP or UDP-Syslog messages) for telemetry received from tags. Therefore, any external applications (such as paging systems, text messaging, enterprise management consoles and so on) relying on northbound notifications in these formats must receive them from an alternate source having visibility to tag telemetry, such as a third-party location client.

Battery telemetry, however, is an exception. In this case, the location appliance will trigger northbound asynchronous notifications based on remaining battery life for tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. These notifications are generated as per the following trigger condition definitions:

- Battery Level is Low—Reported battery life remaining is 30%
- Battery Level is Medium—75% battery remaining > 30%
- Battery Level is Normal—Battery remaining is > 75%

Deploying Tag High-Priority Notifications

Beginning with software Release 4.1 of the Cisco UWN, asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification may transmit high-priority and vendor-specific notifications to the location-aware Cisco UWN. This information is transmitted as part of a tag transmission that is sent on-demand, and is passed from WLAN controllers to the Cisco Wireless Location Appliance using LOCP. Keep in mind that the format of the tag message sent by the tag when a high-priority type event occurs is very similar to the standard tag multicast transmission sent during each tag transmission interval, except that it contains additional information that conveys the nature of the high-priority event.

It is important to note that information contained in the tag notifications received over RF by the WLAN controller is passed (with minimal delay) to the location appliance in the form of LOCP Information Notifications. Thus, for example, when a call button is depressed on an asset tag that is compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification, a LOCP Information Notification is transmitted by the WLAN controller to the location appliance very shortly after the tag notification has been received by the controller's registered access points. Once received by the location appliance, the updated call button status is reflected in the location appliance database (for example, "panic button depressed") and made available to location clients. If archiving of tag historical information has been enabled on the location appliance, tag "emergency" information is archived along with other tag information (shown in [Figure 6-28](#)).

The basic size of a LOCP Information Notification Ethernet frame is approximately 130 bytes. Frame sizes can be larger based on additional information included in the frame, such as tampering information or vendor-specific data. In Release 4.1, LOCP Information Notifications are not aggregated by WLAN controllers. WLAN controllers will transmit a LOCP Information Notification frame to the location appliance for *each* tag high-priority notification received via *each* of its registered access points (including any high-priority notification repetitions).

Expressed mathematically, it can be stated that for each notification event coming from a tag, the total number of LOCP Information Notifications that can be expected to be transmitted from a WLAN controller to the location appliance can be calculated as:

$$LOCP\ Information\ Notifications_{TOTAL} = Detecting\ APs_{TOTAL} * High-Priority\ Notification\ Repetitions_{PER\ CHANNEL}$$

where *High-Priority Notification Repetitions_{PER CHANNEL}* represents the total number of high-priority notifications that are sent by the tag on a single RF channel. Note that the number of high-priority notification repetitions per channel should not be confused with the standard setting for tag message notifications per channel, which applies to tag transmissions that are sent periodically based on the expiration of a tag transmission interval. It should also be noted that this calculation yields the maximum possible value for *LOCP Information Notifications_{TOTAL}* as it assumes that all notification repetitions coming from the tag are successfully detected by all access points included within *Detecting APs_{TOTAL}* and none are dropped due to interference, contention or other RF anomalies.

Using our formula, we can calculate the expected number of LOCP Information Notifications that will be generated if the call button is depressed once on an asset tag compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification within the following Release 4.1 environment:

- Two WLAN controllers
- Three access points registered to each controller, for a total of six detecting access points.
- Tags send one notification for each call button depression on each of channels 1, 6 and 11

Substituting this information into the aforementioned equation, we see that $6 * 1$ or 6 total LOCP Information Notifications will be transmitted from the WLAN controllers to the location appliance in this example. Note that although both WLAN controllers will be sources of LOCP Notifications in this example, the number of WLAN controllers present in the environment has no bearing on the number of LOCP Notifications that will be sent to the location appliance. We could have substituted three WLAN controllers with two access points registered to each in this example, and the calculated value for *LOCP*

*Information Notifications*_{TOTAL} would have been the same. It is the number of access points that detect the tag multicast transmissions bearing the high-priority notification information sent that is pertinent to the number of LOCP Notifications that will be generated from controllers to the location appliance.

To ensure proper LOCP operation between the location appliance and any WLAN controllers defined to it, always ensure that port 16113 is not blocked by any firewalls or other security devices.

Configuring Tags for Telemetry and Notifications

While the support of tag telemetry and notifications are basic components of the Cisco Compatible Extensions for Wi-Fi Tags specification, each tag vendor uses their GUI or CLI-based tag software to enable, disable or otherwise customize precisely how these features are supported in their products. While a limited amount of AeroScout tag configuration information has been already provided in prior sections of this document, more comprehensive information specifically relating to the configuration of external telemetry sensors and asset tags is available from asset tag vendors, but is beyond the scope of this document.

Readers seeking such information are directed to the following sources of information:

- *AeroScout T2 Tag User Guide*
- *AeroScout Tag Manager User Guide version 3.0*
- <http://www.aeroscout.com> or your AeroScout account and technical support team

For asset tags from other vendors that are compatible with the Cisco Compatible Extensions for Wi-Fi Tags specification, it is recommended to contact those vendors directly. These would include:

- InnerWireless (formerly PanGo Networks) <http://www.innerwireless.com>
- WhereNet <http://www.wherenet.com>

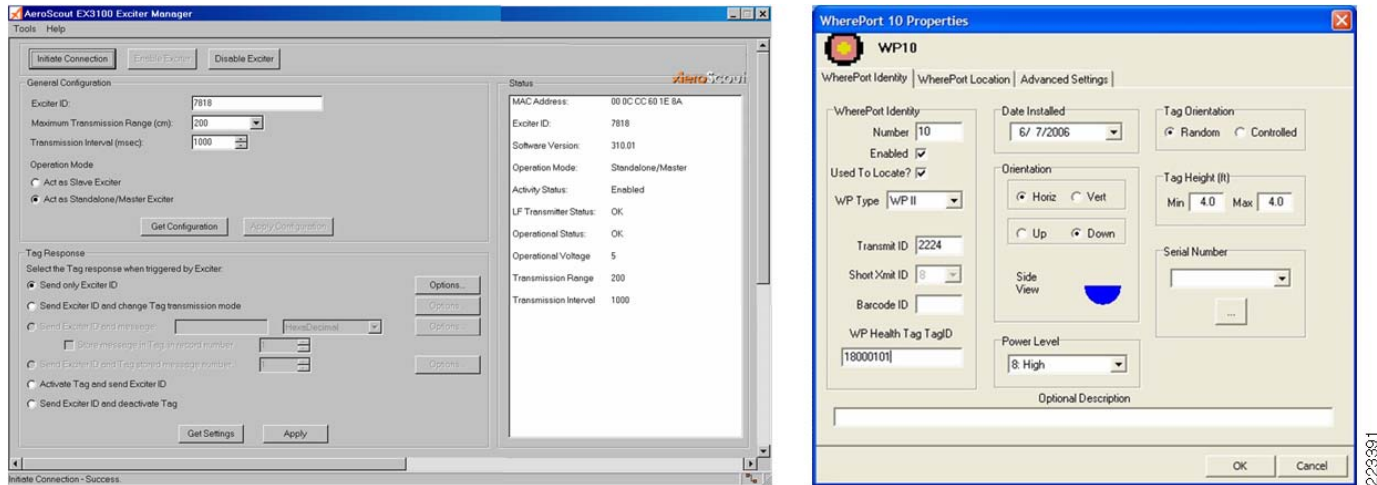
Chokepoint Considerations

Configuring Chokepoint Triggers

In order to use chokepoint triggers with the Cisco UWN, they must be properly configured using the appropriate vendor-supplied software utility, defined to WCS, placed on floor maps and synchronized as part of an updated network design to the location appliance. After all of this is complete, the location appliance will be able to recognize that asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification have been stimulated by a particular chokepoint trigger MAC address and proceed to localize the asset tag. Location clients may then display the asset tag's location at the chokepoint icon associated with the chokepoint trigger's MAC address.

Various chokepoint trigger specific parameters such as transmission range, IP address, transmission interval, transmission repetitions and so on are set using vendor-specific utilities. For non IP-addressable AeroScout EX-3100 series Exciters, the AeroScout Exciter Manager standalone software utility must be used (shown on the left in [Figure 6-29](#)). For WhereNet WherePort chokepoint triggers, the WhereNet System Builder (shown on the right in [Figure 6-29](#)) and the WhereNet WhereWand are used.

Figure 6-29 Vendor-Specific Configuration Utilities



Note that each vendor maintains their set of software tools necessary for configuration of their chokepoint triggers. These software configuration tools are not interoperable between vendors (for example, AeroScout software configuration tools cannot be used to configure WhereNet chokepoint triggers or vice-versa).

In general, the individual configuration of each vendor's chokepoint trigger device is beyond the scope of this white paper. This document does, however, attempt to shed light on specific chokepoint trigger configuration parameters that are of particular significance in solving design challenges. As necessary, the topical sections of this document make reference to such parameters as necessary. However, complete and detailed configuration information relating to the specific configuration of each vendor's chokepoint trigger can be found in the appropriate vendor's documentation:

Available from AeroScout Corporation:

- AeroScout EX-3100 Exciters:
 - *AeroScout Exciter EX-3100 User Manual*
 - *AeroScout EX-3100 Exciter Manager User's Manual*
- AeroScout EX-3200 Exciters:
 - *AeroScout EX-3200 User Guide*
- AeroScout EX-2000 Exciters:
 - *AeroScout Exciter EX-2000 User Guide*

The following reference manuals are recommended for configuration of AeroScout EX-2000 and EX-3200 Exciters, using either the AeroScout System Manager or the AeroScout Network Exciter Manager (ANEM). The AeroScout Network Exciter Manager is a standalone Exciter software configuration utility specifically designed for users of AeroScout Exciters and the Cisco UWN.

- *AeroScout Engine Version 3.2 User's Guide*
- *AeroScout Network Exciter Manager (ANEM) User's Guide*

Technical documentation for WhereNet WherePort chokepoint triggers and the necessary software and hardware for configuration of WherePorts is available from WhereNet Corporation (<http://www.wherenet.com>) via your WhereNet account representative.

Defining Chokepoint Triggers to the Cisco UWN

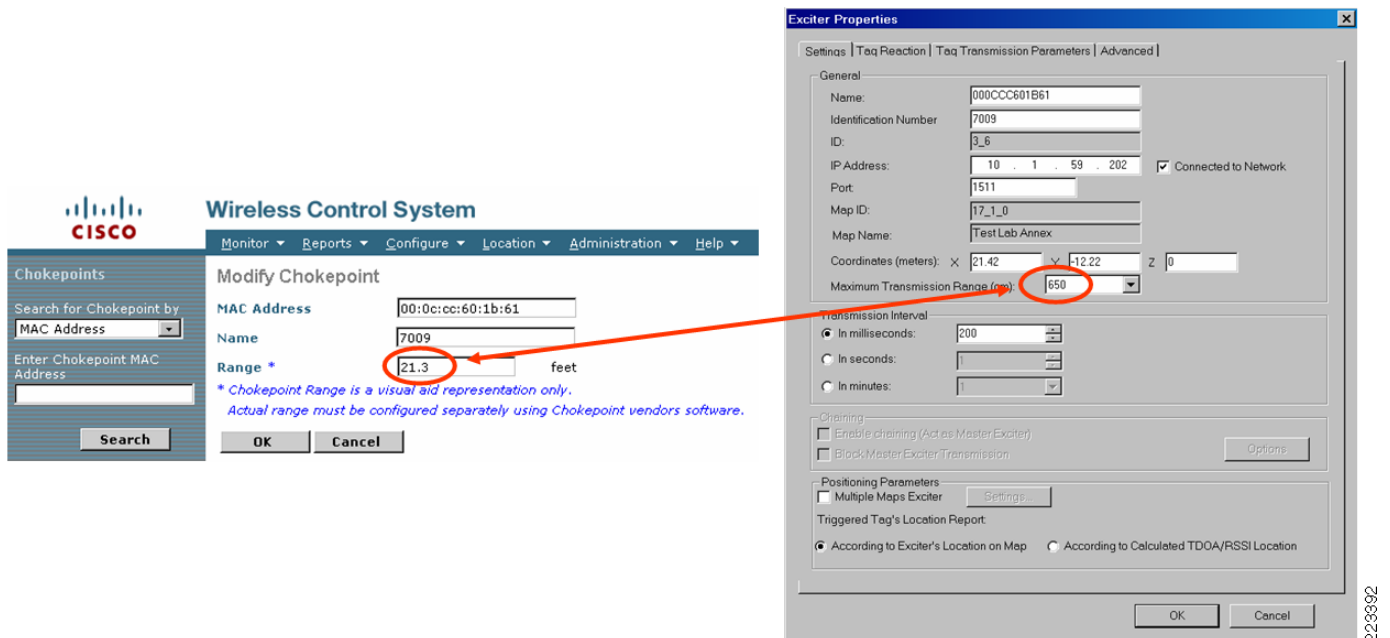
As mentioned earlier, after chokepoint triggers have been individually configured using the configuration tools supplied by the vendor, they must be defined to WCS, placed on appropriate floor maps and synchronized with the location appliance as part of an updated network design. Only then can they be used to track asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification.

Complete step-by-step guidance regarding how to define compatible chokepoint triggers to WCS and the location appliance can be found at the following location:

http://www.cisco.com/en/US/products/ps6386/products_configuration_guide_chapter09186a008082d7d2.html#wp1058654.

When defining chokepoint triggers, it should be noted that the range is specified in both the vendor's configuration program as well as in WCS (shown in Figure 6-30). However, it is the range configuration parameter specified in the vendor's configuration program that actually sets the transmission range of the chokepoint trigger, not the range setting in WCS. The value that is specified for the range of the chokepoint trigger in WCS simply sets the size of the gray concentric rings that appear surrounding each chokepoint icon on WCS floor maps. These concentric rings are visual aids placed simply to serve as a convenient reminder of the range associated with the chokepoint trigger.

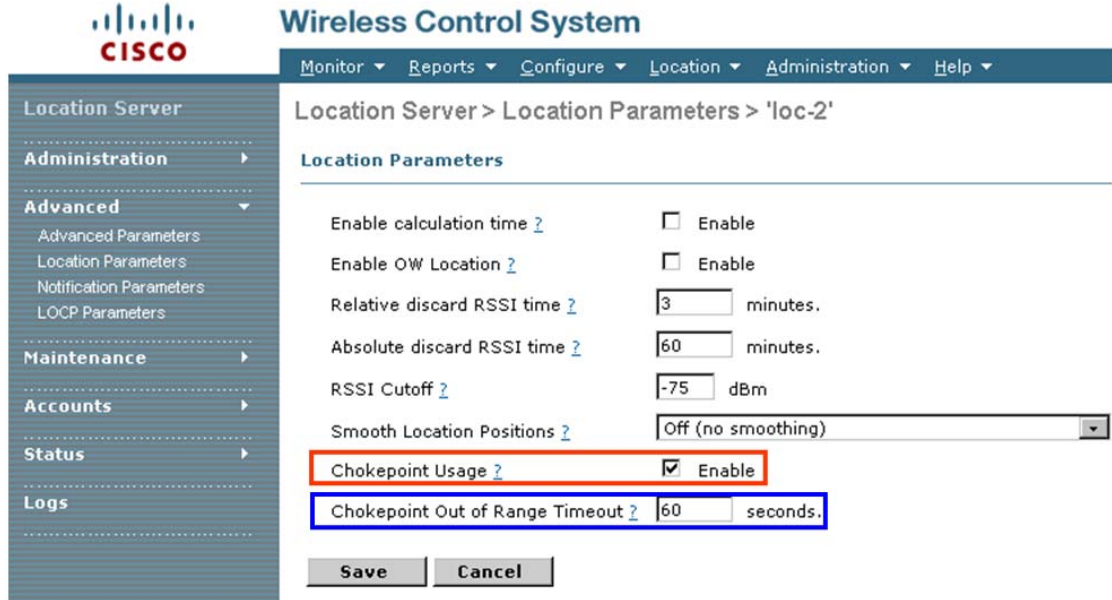
Figure 6-30 WCS and Vendor Range Parameters Compared



Note that these concentric rings do not represent any type of “special” area. For example, when RF Fingerprinting is being used to as the means of localizing tags instead of chokepoint location, tags may be placed by the system anywhere on floor maps (including within these gray concentric rings) if that is the location deemed to be correct by the location appliance.

There are also two additional parameters regarding the use of chokepoints that are found on the Location > Location Servers > Advanced > Location Parameters menu screen, as shown in Figure 6-31:

Figure 6-31 Chokepoint Advanced Location Parameters



- Chokepoint Usage**—This checkbox (shown within the red rectangle in Figure 6-31) must be enabled in order for the location appliance to use chokepoint location techniques to localize tags. This occurs when it receives incoming LOCP Measurement Notifications indicating that a tagged asset has been stimulated by a chokepoint trigger. With regard to the chokepoint capabilities contained with the Cisco location appliance and the location-aware Cisco UWN, these techniques are only used with asset tags that are compliant with the Cisco Compatible Extensions for Wi-Fi Tag specification. If this parameter is disabled, the appliance will use the standard mechanism of RSSI based RF Fingerprinting to calculate tag location at all times.
- Chokepoint Out of Range Timeout**—This parameter (shown within the blue rectangle in Figure 6-31) specifies the timer used to age the last “in-range” report received from for an asset tag that is being localized using chokepoint location techniques. It assures that any tags no longer transmitting frames indicating they are within range of a chokepoint trigger are removed from that chokepoint in the active location database, once the Chokepoint Out of Range Timeout has expired. These tags are assumed to have left the chokepoint and are reverted back to being localized using standard RF Fingerprinting techniques.

Chokepoint Trigger Traffic Considerations

Beginning with Cisco UWN software Release 4.1, tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification can use a consistent method to inform the UWN that they are within (or have left) the proximity of a chokepoint trigger. Once received by access points and forwarded to registered controllers, this information is passed to the location appliance using LOCP Measurement Notifications, which have already been described in [Cisco Location Control Protocol \(LOCP\)](#), page 3-36.

The length of each 802.11 tag multicast frame transmitted in response to stimulation received from a chokepoint trigger is approximately 63 bytes, which includes only a single chokepoint MAC address and does not include any historical chokepoint information. The length of the frame could increase due to the inclusion of a historical list of chokepoints traversed, or it may be larger than 63 bytes due to vendor-specific information that may be included in the frame. For example, during lab testing with

AeroScout T2 tags, it was observed that the typical size of the tag multicast frame emitted when in proximity of a chokepoint trigger is approximately 71 bytes, slightly larger than the multicast frame transmitted by these same tags during routine periodic transmissions (56 bytes). This 71-byte length is greater than the expected 63 bytes, and upon further examination it is discovered that eight additional bytes of vendor-specific information is included.

The Cisco Compatible Extensions for Wi-Fi Tags specification also allows asset tags to communicate historical information about the chokepoints they traverse to the Cisco UWN. This could increase the size of the frame by approximately 10 bytes per chokepoint trigger encountered depending on the number of historical entries maintained. The basic size of a LOCP Measurement Notification Ethernet frame is approximately 160 bytes. Frame sizes may be larger based on additional information included in the frame, such as historical chokepoint information.

In software Release 4.1, LOCP Measurement Notifications are not aggregated by WLAN controllers. WLAN controllers will transmit a LOCP Measurement Notification frame to the location appliance for each incoming tag multicast transmission, received by each of its registered access points, that indicates that the tag has been successfully stimulated by a chokepoint trigger. Therefore, the number of LOCP Measurement Notifications generated by one or more WLAN controllers for a single tag transmitting multicast frames indicating that the tag has been stimulated by a chokepoint trigger, is dependent upon:

- the number of registered access points that are within range of the tag and that have detected the tag's chokepoint-related transmissions.
- the number of times the tag will transmit a multicast frame on each configured 802.11 channel in response to chokepoint trigger stimulation.

This can be expressed mathematically as:

$$LOCP\ Measurement\ Notifications_{CHOKEPOINT} = Detecting\ APs_{TOTAL} * 802.11\ Repetitions_{PER\ CHANNEL}$$

Note the following considerations:

1. This calculation yields the number of *LOCP Measurement Notifications* that result from a single tag reacting to a single chokepoint stimulation event.
2. Chokepoint triggers by default transmit multiple stimulation packets over their magnetic signaling medium. This could result in multiple stimulation events, which is highly dependent on the amount of time spent within the chokepoint stimulation zone and other factors.
3. This calculation yields a maximized value for *LOCP Measurement Notifications* as it assumes that all frames transmitted by the tag are successfully detected by the number of access points specified in *Detecting APs_{TOTAL}* (none are dropped due to interference, contention or other RF anomalies).

In the majority of cases:

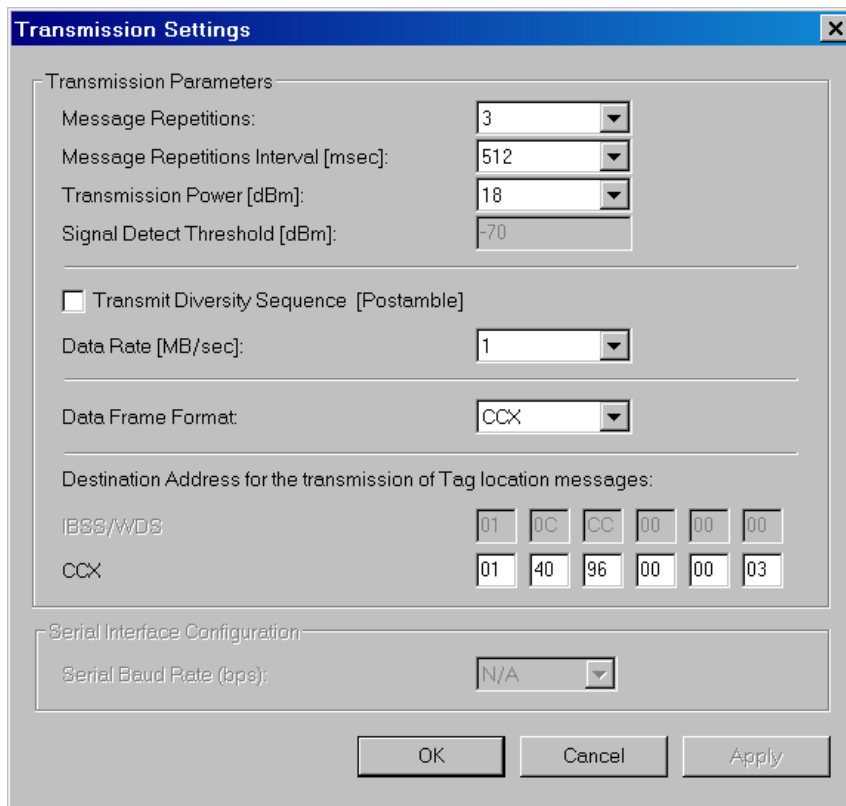
- Chokepoints are deployed in areas where the surrounding access point spacing meets the requirements discussed in [Access Point Placement, page 5-5](#).
- Access points and tags are configured to operate on the non-overlapping 2.4 GHz channels (channels 1 (2412 GHz), 6 (2437 GHz) and 11 (2462 GHz) in the Americas, for example).

Most chokepoint triggers assume a default value of one for the number of times they repeat, per channel, tag multicast transmissions indicating that the tag has been successfully stimulated by a chokepoint trigger. Only a single tag multicast transmission frame containing the stimulating chokepoint trigger's MAC address need be received in order to result in the generation of a *LOCP Measurement Notification*. Because of this, the default value for the number of times these chokepoint-related tag transmissions are repeated per channel is usually sufficient, especially since this tag transmission will typically be repeated across three 2.4 GHz channels, resulting in more than one access point receiving the tag transmission, even without increasing the repetition count. However, in some cases where interference or congestion may be extremely high, it may make sense to increase the repetition count slightly. In other cases involving tagged assets traversing through chokepoint areas at high speed or at fringe distances from

chokepoint triggers, this parameter can be used to increase the likelihood of reliable stimulation (see [Appendix A, “Chokepoint Transmission Interval Analysis”](#) for more details). In all cases, however, such increases should be done judiciously given the ability of this parameter to affect the amount of traffic added per stimulated tag in large tag environments.


It should be noted that the repetition count that applies to the tag multicast frames sent in response to chokepoint stimulation is usually managed independently of the repetition count for other tag events such as telemetry, high-priority notifications or periodic tag transmissions sent as a result of the tag's configured transmission interval. When configuring tags and chokepoint triggers, it is important to maintain this distinction. For example, with AeroScout tags the repetition count that applies to the 802.11 frames sent by a tag in response to a chokepoint stimulation event is known as the “Tag Repetition of an Exciter event” parameter. It is configured on a per-Exciter basis using the AeroScout System Manager, Exciter Manager or ANEM utility. In contrast, the tag repetition parameter used for non-Exciter related events is known as the Message Repetitions transmission parameter. It is set on a per-tag basis using the Transmission Settings panel of the AeroScout Tag Manager, as shown in [Figure 6-32](#).

Figure 6-32 Transmission Settings Panel in AeroScout Tag Manager (not used for Exciter Events)



As mentioned earlier, the length of each 802.11 multicast tag frame transmitted by a tag in response to chokepoint stimulation is approximately 63 bytes, which includes only a single chokepoint MAC address and does not include any historical chokepoint information. We mentioned earlier that it was observed during testing that the 802.11 multicast frame transmitted by an AeroScout T2 tag also contains eight bytes of vendor-specific information. [Figure 6-33](#) illustrates this, with the mandatory chokepoint information contained within the yellow rectangle and the additional vendor-specific information contained within the blue rectangle.

Figure 6-33 Vendor-Specific Information Included in Tag Chokepoint Transmission



Although both the standard chokepoint information as well as the optional vendor-specific information travels from the tag to the access point contained within the same tag multicast frame, in software Release 4.1 the WLAN controller parses this into two separate LOCP notifications:

- An LOCP Measurement Notification containing the chokepoint group information that is 160 bytes in length.
- An LOCP Information Notification containing the vendor-specific information that is 138 bytes in length.

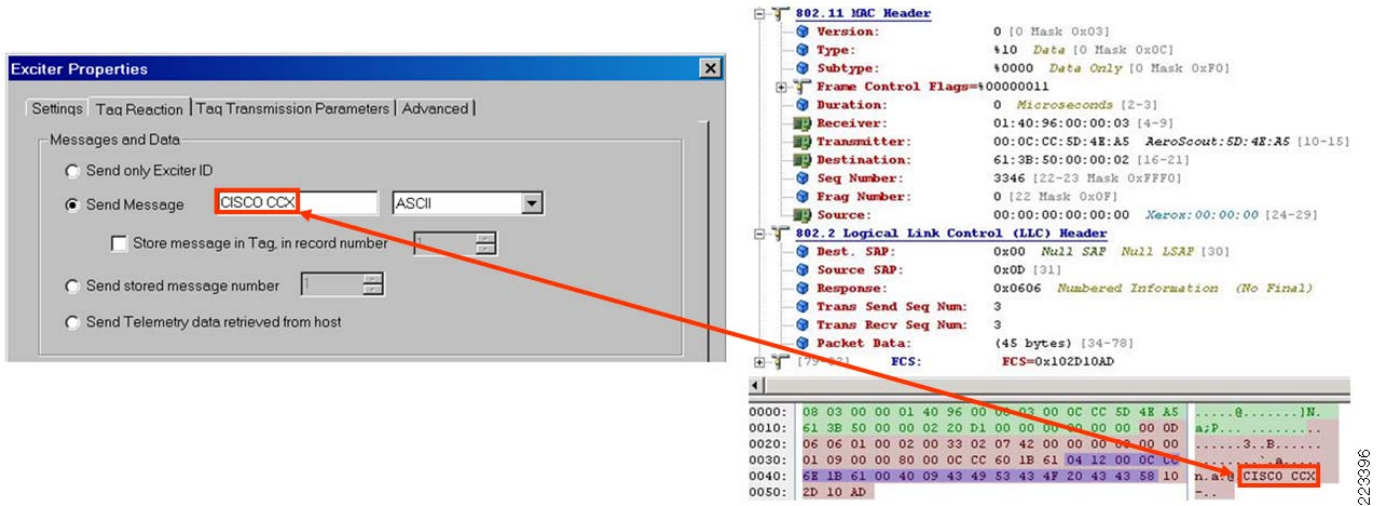
The precise composition of the vendor-specific fields varies depending on the chokepoint and tag vendor. For example, AeroScout allows for additional message information to be appended to the Exciter ID via the Tag Reaction tab of the Exciter Properties menu in the AeroScout System Manager and AeroScout Network Exciter Manager. These capabilities are also available via the Exciter Manager utility for users of the AeroScout EX-3100 Exciter.

Vendor specific information can be:

- Directly entered and saved on a per-Exciter basis.
- Saved to tag memory for later reuse.
- Consist of one of ten preconfigured messages programmed into tags.
- Emanate from a host attached to the tag.

Figure 6-34 illustrates the use of this capability for an AeroScout EX-2000 Exciter and AeroScout tags. In this figure the Exciter instructs the tag to append vendor-specific information in addition to the vendor-specific Exciter ID to each tag transmission frame sent as a result of stimulation received from the Exciter.

Figure 6-34 AeroScout Vendor-Specific Information Options



In Figure 6-34 we see the message “CISCO CCX” being defined to the Exciter as well as the complete 83 byte message transmitted by the tag when stimulated by the Exciter. This 83-byte message includes the standard information regarding the MAC address of the stimulating Exciter as well as the vendor specific information. Note that the text defined to the Exciter in the AeroScout System Manager is seen transmitted by the tag at offset x0046 in the trace (you can see the ASCII text “CISCO CCX” shown at the right in Figure 6-34). Every access point receiving this information will forward it to their registered controller where a 160-byte LOCP Measurement Notification as well as a 148-byte LOCP Information Notification will be sent to the location appliance. Although this information was hard-coded at the Exciter, the Exciter could have just as easily instructed the tag to instead include telemetry data that it retrieved from the asset (host) that it is attached to, such as embedded sensor data.

Keep in mind that results of our test observations obviously are, in this case, AeroScout specific, as other vendors may or may not opt to allow the inclusion of vendor-specific information to the same degree.



CHAPTER 7

Caveats

This chapter provides the caveats discovered in lab testing.

CSCse14724—Degraded Location Accuracy with Monitor Mode APs

Degraded accuracy has been observed in lab testing of monitor mode access points when compared to local mode.

The use of Monitor Mode in location aware designs with software Release 4.1 is not recommended at this time.

CSCsh88795—CCX S36 Beacon Measurement Request Dual-Band Support

CCX S36 Beacon Request includes channels from the same band as association but not the other band. This can affect the reliability of performing simultaneous calibration data collection on both bands when using dual-band clients. The band currently associated will typically calibrate reliably, whereas the other band does not experience the same degree of reliable probe-request generation that is brought about by the use of unicast Radio Measurement Requests.

Workaround

It is recommended that calibration data collection be performed for each band individually at this time, even when using dual-band clients. To accomplish this, use either of the following alternatives:

1. Perform the calibration data collection on each band individually using a single laptop equipped with a dual-band client adapter compatible with the Cisco Compatible Extensions specification for WLAN devices specification at version 2 or higher, and capable of recognizing and responding to S36 unicast radio measurement requests. An example of such a client is the Cisco Aironet 802.11a/b/g Wireless CardBus Adapter (AIR-CB21AG). For example, proceed to disable the 5 GHz band and complete the data collection using the 2.4 GHz band only. Then, disable the 2.4 GHz band and enable the 5 GHz band, and proceed to repeat the data collection using the 5 GHz band only.
2. Perform the calibration data collection using two operators and two independent laptops. Each laptop should be equipped with a dual band client adapter compatible with the Cisco Compatible Extensions specification for WLAN devices specification at version 2 or higher, and capable of recognizing and responding to S36 unicast radio measurement requests. An example of such a client

is the Cisco AIR-CB21AG. Each laptop should be associated to the infrastructure using a different band. The two calibration data collection operators may function independently; there is no need for them to visit each data point at the same time, or to even visit the same data points. In this way, a complete calibration data collection can be performed across both bands in half the time as option #1 above.

CSCsi95122—WCS Does Not Dispatch Northbound Emails for Location Notifications

WCS does not send email notifications for any location notification alarm categories. Alarms for location notifications appear on the alarm console, however email notifications do not get transmitted.

Workaround

Use email northbound notifications present in the Location Appliance instead of WCS with software Release 4.1.

For additional caveats than those discussed above, refer to the following documents:

- Release Notes for Cisco Wireless Location Appliance 3.0—http://www.cisco.com/en/US/products/ps6386/prod_release_notes_list.html
- Release Notes for the Cisco Wireless Control System (WCS) 4.1—http://www.cisco.com/en/US/products/ps6305/prod_release_notes_list.html
- Release Note for Cisco WLAN Controllers and Lightweight Access Points 4.1—http://www.cisco.com/en/US/products/ps6366/prod_release_notes_list.html
- Release Notes for Cisco Aironet Access Points for Cisco IOS Release 12.3(11)JA1—http://www.cisco.com/en/US/products/hw/wireless/ps430/prod_release_notes_list.html
- Cisco Bug Toolkit—<http://www.cisco.com/cgi-bin/Support/Bugtool/home.pl>



APPENDIX **A**

Determining Approximate Roots using Maxima

In the circle-circle intersection equations below:

- d represents the inter-access point distance in feet
- Obg represents the percentage of overlap desired for 802.11bg
- Oa represents the percentage of overlap desired for 802.11a
- R represents the cell radius in feet

Note that zero overlap occurs when the distance between the centers of the two circles is equal to twice the radius ($d = 2R$). This relationship becomes invalid if d is allowed to exceed $2R$, as an area of intersection would be impossible to calculate. This is why the *find_root* function is limited to the closed interval from $d/2$ to d .

```
wxMaxima 0.7.3a http://wxmaxima.sourceforge.net
Maxima 5.13.0 http://maxima.sourceforge.net
Using Lisp GNU Common Lisp (GCL) GCL 2.6.8 (aka GCL)
Distributed under the GNU Public License. See the file COPYING.
Dedicated to the memory of William Schelter.
(%i1) d:45.9;
(%o1) 45.9
(%i2) Obg:0.10;
(%o2) 0.1
(%i3) Oa:0.15;
(%o3) 0.15
(%i4) find_root(Obg*pi*R^2=2*R^2*acos(d/(2*R))-(1/2*d*(sqrt(4*R^2-d^2))),R,d/2,d);
(%o4) 28.49573663945017
(%i5) find_root(Oa*pi*R^2=2*R^2*acos(d/(2*R))-(1/2*d*(sqrt(4*R^2-d^2))),R,d/2,d);
(%o5) 30.87736860938116
```



Note

Maxima is not produced, marketed, sold, or supported by Cisco. Maxima is a publically available computer algebra system (CAS) that has been released under the GNU Public License. Further details regarding Maxima, its capabilities, and its use (including downloads for various operating systems) can be found at the following URL <http://maxima.sourceforge.net>.



APPENDIX **B**

Verifying Detection of Asset Tags in WLAN Controllers

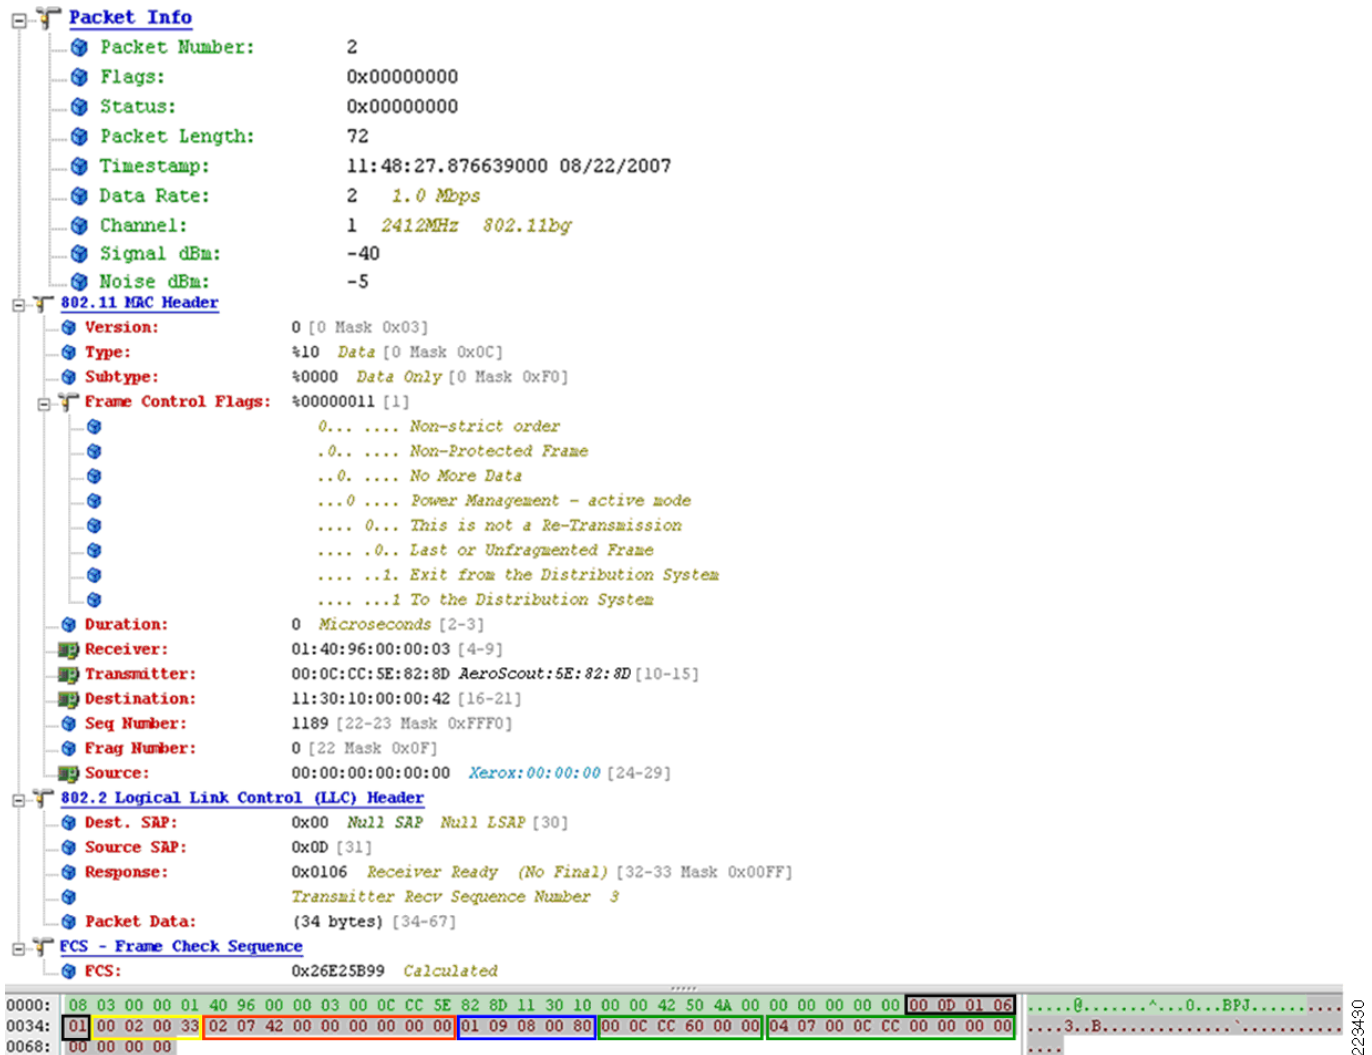
Asset Tags Detection

The protocol analyzer trace in [Figure B-1](#) provides important information with regard to how asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification are recognized and distinguished from other tracked devices in the network. In this example, we use an AeroScout T2 asset tag with firmware version 4.33. Assets tags that are supplied by other vendors that are also compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification can be expected to be recognized by the Cisco UWN in a similar fashion.

[Figure B-1](#) depicts the layer two multicast frame that is transmitted at the expiration of every tag transmission interval for an AeroScout T2 asset tag configured for a basic set of operational parameters. In [Figure B-1](#), the tag configuration includes:

- Periodic 60 second tag transmission interval across three channels (1, 6, 11)
- Chokepoint out of range indication (indicated by the blue rectangle in [Figure B-1](#))
- Onboard motion and temperature detection sensors disabled

Figure B-1 RF Protocol Analysis of Tag Multicast Frame (Cisco Compatible Extensions for Wi-Fi Tag Compliant)



This asset tag shown in Figure B-1 is not configured to transmit external sensor telemetry. In addition, the RF frame also includes the following information:

- Five byte Cisco Compatible Extensions for Wi-Fi Tags header (black rectangle)
- Tag product type identification (yellow rectangle)
- Optional battery telemetry (red rectangle)
- Optional vendor specific information (green rectangles)

The length of the multicast frames is dependent upon the tag's configuration and the optional features supported by the tag and tag vendor. In this case, the length of the multicast frame shown in Figure B-1 is 72 bytes. If additional features such as on-board temperature sensing were enabled, or if the tag were transmitting a multicast message due to stimulation from a chokepoint trigger, the frame length would be greater. For example, a typical length for tag multicasts transmitted as a result of stimulation from a chokepoint trigger is 88 bytes. The added length in this case comes primarily from the inclusion of the stimulating chokepoint trigger's MAC address and additional vendor-specific information.

The AeroScout T2 tag initiates Clear Channel Assessment (CCA) for 100 microseconds. If the channel is clear, it then multicasts its payload at 1 Mbps. These frames are sent at 1 Mbps with the To Distribution System (ToDS) and Exit From Distribution System (FromDS) bits in the 802.11 MAC header both set to “1”. Note that the Wireless Distribution System (WDS) four-address frame format is being used, indicated by the presence of the receiver and transmitter addresses in [Figure B-1](#).

The transmitter address will always indicate the MAC address of the asset tag responsible for transmitting the frame, whereas the receiver address is a multicast address used by all asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification, regardless of vendor origin. The destination and source addresses shown within the 802.11 MAC header are not used by the Cisco UWN for asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification. These are typically set to all zeroes, although vendor-specific usage of the destination address field by tag vendors is possible, as we see with the AeroScout T2 tag shown in [Figure B-1](#).

After the frame shown in [Figure B-1](#) is received by access points, it will be transmitted to the controller(s) to which these access points are registered using the LWAPP protocol, as shown in [Figure B-2](#). Here we see the IP source address associated with the receiving access point, and the IP destination address associated with the AP Manager interface of the controller to which the receiving access point is registered. When comparing the two figures, notice in [Figure B-2](#) that Cisco Aironet access points make two modifications to the frame information prior to dispatching to the controller via LWAPP:

- It copies the access point’s base radio MAC address (base BSSID) to the receiver address field in the encapsulated 802.11 header.
- It copies the CCX multicast address of 01:40:96:00:00:03 to the destination address field in the encapsulated 802.11 header.

Figure B-2 LWAPP Capture of Tag Multicast Frame (Cisco Compatible Extensions for Wi-Fi Tag Compliant)

No.	Time	Source	Destination	SrcPort	DstPort	Protocol	Bytes	Info
6385	2232.393942	10.1.95.252	10.1.92.19	54420	12222	LLC	116	S, func=RR, N(R)=3; DSAP
<pre> Frame 6385 (116 bytes on wire, 116 bytes captured) Ethernet II, Src: Cisco_ed:49:44 (00:14:1c:ed:49:44), Dst: Airespac_40:98:03 (00:0b:85:40:98:03) Internet Protocol, Src: 10.1.95.252 (10.1.95.252), Dst: 10.1.92.19 (10.1.92.19) User Datagram Protocol, Src Port: 54420 (54420), Dst Port: 12222 (12222) Source port: 54420 (54420) Destination port: 12222 (12222) Length: 82 Checksum: 0x828c [correct] LWAPP Encapsulated Packet Version: 0 slotId: 0 0.. = Type: Encapsulated 80211 0. = Fragment: Set 0 = Fragment Type: Set Fragment Id: 0x00 Length: 68 RSSI: 0xc9 SNR: 0x2c IEEE 802.11 Type/Subtype: Data (0x20) Frame Control: 0x0308 (Swapped) Version: 0 Type: Data frame (2) Subtype: 0 Flags: 0x3 Duration: 0 Receiver address: Cisco_59:41:f0 (00:14:1b:59:41:f0) Transmitter address: Aeroscou_5e:82:8d (00:0c:cc:5e:82:8d) Destination address: 01:40:96:00:00:03 (01:40:96:00:00:03) Fragment number: 0 Sequence number: 1189 Source address: 00:00:00_00:00:00 (00:00:00:00:00:00) Logical-Link Control Data (34 bytes) 0000 00 0b 85 40 98 03 00 14 1c ed 49 44 08 00 45 00 . . @ ID . . E . 0010 00 66 00 1c 00 00 ff 11 eb 59 0a 01 5f fc 0a 01 . f Y 0020 5c 13 d4 94 2f be 00 52 82 8c 00 00 00 44 c9 2c \ . . / . R D . 0030 03 08 00 00 00 14 1b 59 41 f0 00 0c cc 5e 82 8d Y A A . 0040 01 40 96 00 00 03 50 4a 00 00 00 00 00 00 00 0d . @ P J 0050 01 06 01 00 02 00 33 02 07 42 00 00 00 00 00 00 3 . B 0060 01 09 08 00 80 00 0c cc 60 00 00 04 07 00 0c cc 0070 00 00 00 00 </pre>								

When the tag multicast address is recognized by the controller, the identity and type of sender is established via the payload information contained in the frame. Depending on the type of information contained within the tag payload, it will be passed to the location appliance using either traditional SNMP poll responses or the new LOCP introduced in software Release 4.1.

Note the sequence number (1189) and fragment fields (0) that appear in both the RF as well as the LWAPP frame analysis. This is an important piece of information that can be very useful when matching packets that flow into access points via 802.11 and out of them via LWAPP. The sequence number for a particular tag frame indicates the number of the tag message and is assigned from a single modulo 4096 counter starting at zero, and is incremented by 1 for each tag message cycle. The fragment number specifies the specific frame within a burst of frames transmitted on a single channel. The fragment number should always start from zero even if the burst length is zero. For a packet burst, the fragment number should be set to n where n is the packet index within the burst starting from 0. For example, if a tag is configured to transmit a burst of length 5, then the fragment number would start at 0 for the first packet in the burst, 1 for the second packet and so on up to 4 for the last packet in the burst.

For example, assume that an asset tag is configured to send a burst length of 3 packets for each of channels 1, 6, and 11. In the case of an AeroScout asset tag, the burst length is configured by using the tag message repetitions transmission parameter. The expected fragment and sequence numbers would be as shown in [Table B-1](#).

Table B-1 Packet Fragment and Sequence Numbers

Packet Instance	Fragment Number	Sequence Number	Channel
0	0	0	1
1	1	0	1
2	2	0	1
3	0	0	6
4	1	0	6
5	2	0	6
6	0	0	11
7	1	0	11
8	2	0	11

Assume a second asset tag is configured to send a single message on each of channels 1, 6, and 11 every 60 seconds. The expected fragment and sequence numbers occurring over the next 120 seconds would be as shown in [Table B-2](#) to [Table B-4](#).

Table B-2 Time+0

Packet Instance	Fragment Number	Sequence Number	Channel
0	0	0	1
1	0	0	6
2	0	0	11

Table B-3 Time+60

Packet Instance	Fragment Number	Sequence Number	Channel
0	0	1	1
1	0	1	6
2	0	1	11

Table B-4 Time+120

Packet Instance	Fragment Number	Sequence Number	Channel
0	0	2	1
1	0	2	6
2	0	2	11

Asset Tags Not Detected

In situations where asset tags are not being detected properly despite configuration of the system in accordance with best practices, re-verification of proper configuration should be performed. It is also recommended that verification of proper asset tag RSSI detection and message forwarding be conducted. The following steps are recommended to accomplish this:

Step 1 Verify if tag is properly detected by WLAN controllers by using the **show rfid summary** command:

```
(Controller) >show rfid summary
```

```
Total Number of RFID : 12
```

RFID ID	VENDOR	Closest AP	RSSI	Time Since Last Heard
00:0c:cc:5d:4c:5e	Aerosct	AP0014.6a1b.41f0	-34	24 seconds ago
00:12:b8:00:20:52	G2	AP001a.a10e.2ffa	-61	16 seconds ago
00:14:7e:00:30:a1	Pango	AP0014.6a1b.41f0	-65	2 seconds ago

If the controller does not detect the tag, use the command **show rfid config** to verify that RFID tag detection has been enabled on the controller.

```
(Controller) >show rfid config
```

```
RFID Tag data Collection..... Enabled
RFID Tag Auto-Timeout..... Disabled
RFID timeout..... 1200 seconds
```

Step 2 If the RFID tag detection is not enabled, enable it using the command shown below. Note that starting with the Cisco UWN software Release 4.1, RFID tag detection is enabled by default.

```
config rfid status enable
```

Step 3 Ensure that the RFID tag timeout is set to a recommended minimum of three times (and a recommended maximum of eight times) the longest tag transmission interval found in the tag population, inclusive of stationary as well as any “in-motion” tag transmission intervals. The valid range of values for this parameter is 60 to 7200 seconds and the default value is 1200 seconds.

For example:

```
(Controller) >config rfid timeout 1200
```

Step 4 Check that the RSSI expiry timeout are set as follows:

```
(Controller) >show advanced location summary
```

```
Advanced Location Summary :
```

```
Algorithm used: Average
Client RSSI expiry timeout: 150 sec, half life: 60 sec
Calibrating Client RSSI expiry timeout: 30 sec, half life: 0 sec
Rogue AP RSSI expiry timeout: 1200 sec, half life: 120 sec
RFID Tag RSSI expiry timeout: 1200 sec, half life: 120 sec
```

If the values are different from default as shown above, set them to default using the following configuration commands:

```
config advanced location expiry {calibrating client | client | rogue-aps | tags }
<seconds>
```

```
config advanced location rssi-half-life {calibrating client | client | rogue-aps | tags}
<seconds>
```

- Step 5** If asset tags are still not detected by the controller using the **show rfid summary** command, enable the following debugs on the controller:

```
debug mac addr <tag mac addr>
debug dot11 rfid enable
00:0c:cc:5e:82:8d Parsing Cisco Tag RFID packet 68
00:0c:cc:5e:82:8d System group 51
00:0c:cc:5e:82:8d Battery group: status 0x42, days 0, age 0
00:0c:cc:5e:82:8d Chokepoint group, option 0x8, power 0, range 128
00:0c:cc:5e:82:8d Vendor group
00:0c:cc:5e:82:8d LOCPBuffer 0x133245ec buffer 0x13324611 msgLen 37 msgId 18 transId
816848706
00:0c:cc:5e:82:8d Notifying LBS of vendor specific data
00:0c:cc:5e:82:8d rfid Aerosct updated by AP 00:14:1B:59:41:F0 (Incoming rssi -47,snr 53),
New saved values rssi -48, snr 49, timestamp 3402186024
```

**Note**

It is recommended that the **debug mac addr** command be used when debugging in this fashion. This will help avoid a flood of debug messages in environments where there are many tags active.

- Step 6** If the **debug** command output indicates that packets from this asset tag are not received by the controller, you may want to continue debugging on the console of an access point that is known to be within range of the asset tag. In order to verify the detection of an RFID tag by an access point, perform the following steps:

- a. Verify whether RFID tag detection has been enabled on the access point and the channel that the access point is currently configured for. This can be done on the access point console using the following command:

```
show controller Dot11Radio 0

<snipped capture>
Current Frequency: 2412 MHz Channel 1
RFID Tag Detection: Enabled
```

If tag detection is found to be disabled on the access point, enable tag detection by issuing the following command on the controller:

```
config rfid status enable
```

Note that some access point commands can also be executed remotely from the controller using the access point remote debugging feature:

```
debug ap enable <Cisco AP>
debug ap command <command> <Cisco AP>
```

- b. If tag detection is enabled and the asset tag is configured to transmit on the channel the access point is configured for, check to see that the access point is forwarding the tag multicast packet to the controller by enabling the following debugs on the access point:

```
debug dot11 Dot11Radio 0 trace print mcast
```

**Note**

In software Release 4.1, in order for the output of the **debug dot11 Dot11Radio 0 trace print mcast** command to be viewed, the command must be entered directly at the access point console. It cannot be entered remotely from the controller.

```
4A1F8FB0 r 1 60/ 38- 0803 000 m014096 5D33CF m61356B 6D00 000000 173
0012 0606 0100 0200 3302 0742 0000 0000 0000 0302 0A03 0109 0000 8000
0CCC 6010 BF04 0800 0CCC 6E10 BF00 0000 019C 917C 00B8 F564 608F FD05 0000
```

- **m014096**—The 3 bytes following the letter “m” represents the first 3 bytes of the multicast address used for asset tags compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification.
- **5D33CF**—This represents the last 3 bytes of the asset tag’s MAC address.

If you would like to view the multicast address and tag MAC address in their unabbreviated format, issue the command **no debug dot11 Dot11Radio 0 print short** command at the access point console.

If information similar to that shown above is seen on the debug output, it indicates that the access point is receiving and forwarding asset tag packets to the controller. If the controller still does not show the asset tag packets being received, use an ethernet protocol analyzer capable of decoding LWAPP encapsulated 802.11 frames (such as WireShark or OmniPeek) on the LWAPP ethernet connection between the access point and controller to verify that the asset tag packets are indeed reaching the controller. The format of these packets should similar to that shown in [Figure B-2](#). If these packets are seen on the protocol analyzer trace and the controller still does not indicate that asset tag packets are successfully received, capture all the details collected so far including the protocol analyzer traces and contact the Cisco Technical Assistance Center for further debugging assistance.

- c. If the tag multicast messages are not seen in the access point debug output, use an RF protocol analyzer such as OmniPeek or WireShark to verify that asset tags are indeed successfully transmitting packets in the format expected on all three 2.4 GHz channels (or the channels that your infrastructure is configured for) as seen in [Figure B-1](#). If the proper frame formats are not seen on the protocol analyzer trace, this should be addressed via the asset tag configuration or by replacing the asset tag if necessary, especially if the asset tag firmware is out of date. If the proper frames are seen on the RF protocol analyzer, attempt to reset tag detection in the controller by issuing the following commands:

```
config rfid status disable
config rfid status enable
```

If the issue continues to persist despite these suggestions, it is recommended that you capture all the details collected so far including the protocol analyzer traces and contact the Cisco Technical Assistance Center for further debugging assistance.

Verifying Asset Tag Telemetry and Events

In order to verify that WLAN controllers are detecting asset tag telemetry and high-priority events, and forwarding that information to the location appliance using LOCP, the following procedure may be used:

-
- Step 1** Make sure that the asset tag is detected by the WLAN controller using the procedure outlined in the previous section.
 - Step 2** Verify that the telemetry or high-priority “emergency” event information you are concerned with has been recorded in the RFID database on the controller:

```
show rfid detail <tag mac addr>
```

For example, the following output snippet illustrates that indication of a panic button being depressed on an asset tag has been received, along with vendor specific data pertaining to the event:

```
<snip>
Telemetry Group
=====
Motion Probability..... No Motion

!! EMERGENCY !!
=====
Reason..... Panic Button

Vendor Specific
=====
Group Length..... 8
Vendor OUI:..... 0x0 0xc 0xcc
Vendor Data: 0x6e 0x13 0xa3 0x0 0x0
```

- Step 3** If the notification is not seen above in the controller's RFID database, enable the following debugs on the controller to validate that it is receiving the notifications.:

```
debug mac addr <tag mac addr>
debug dot11 rfid enable
```

```
00:0c:cc:5e:82:8d Parsing Cisco Tag RFID packet 62
00:0c:cc:5e:82:8d System group 51
00:0c:cc:5e:82:8d Battery group: status 0x42, days 0, age 0
00:0c:cc:5e:82:8d Chokepoint group, option 0x8, power 0, range 128
00:0c:cc:5e:82:8d Emergency group
00:0c:cc:5e:82:8d LOCP Buffer 0x133245ec buffer 0x1332460a msgLen 30 msgId 18 transId
808989330
00:0c:cc:5e:82:8d Notifying LBS of emergency
00:0c:cc:5e:82:8d rfid Aerosct updated by AP 00:14:1B:59:41:F0 (Incoming rssi -49,snr 50),
New saved values rssi -48, snr 49, timestamp 3444175997
```

- Step 4** If output similar to that above is not seen, use an RF protocol analyzer to capture the packets being transmitted by the tag during the send of telemetry or high priority events to verify that data is being transmitted in the correct format. You will need the assistance of the Cisco Technical Assistance Center to confirm this. If the packets that the tag transmits over the air are deemed by the Cisco TAC not to be valid, verify that the level of firmware being used in the asset tag is compliant with the Cisco Compatible Extensions for Wi-Fi Tags specification and that it supports the telemetry or high-priority functions desired.

- Step 5** If the telemetry or high-priority events are seen in the controller's RFID database, check to see if they are being sent to the location appliance (look for **Notifying LBS of emergency**) in the output above. Issue the following command to verify that the LOCP connection between the controller and the location appliance is up and functioning.

```
(Cisco Controller) >show LOCP status
```

LocServer IP	TxEchoResp	RxEchoReq	TxData	RxData
10.1.56.21	5300	5300	83597	441

Normally, if the LOCP connection is up and running properly, you should see the echo counts regularly increment (based on the settings of the echo interval in the location appliance). In addition, as emergencies and telemetry events occur that require transport via LOCP, you will see the data fields

increment as well. If a controller is rebooted and repeatedly fails to establish a connection to the location appliance, you will fail to see an IP address listed for the location appliance, and the Tx / Rx counts will be blank.

If the TxData fields fail to increment in spite of known emergencies and telemetry data being sent by tags, verify that the LOCP send to the location appliance is successful using the following debug command:

```
debug LOCP event enable
```

The output should look similar to the following:

```
LOCP TX message  
Sending LOCP_APP_INFO_NOTIF_MSG to LocServer 0  
Tx OK
```

If messages are received indicating that there are LOCP failures, contact the Cisco Technical Assistance Center for further troubleshooting assistance.
