



Wireless Considerations in Healthcare Environments

Version 1.0

May 14, 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

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.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site 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. (0304R)

Wireless Considerations in Healthcare Environments

Copyright © 2008 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1**Wireless LAN Networks in Hospital Environments 1-1**

CHAPTER 2**RF Design Considerations 2-1**

Overview 2-1

Capacity Planning for High Bandwidth Areas 2-4

Channel Planning for High Density Areas 2-6

Designing Access Point RF Coverage to Meet the Needs of Wireless Applications 2-10

Design Considerations for the Location Services and Voice in the Same Site 2-12

Design Considerations for 802.11b/g/n and 802.11a/n in the Same Coverage Area 2-16

Determining the Cell Edge of the Clients 2-17

Using the Wireless Control System to Estimate the Number of Access Points to Support Applications 2-19

Using the Wireless Control System to Maintain High Quality VoWLAN Performance 2-21

Establishing Baselines for Measuring and Managing WLAN Performance 2-26

Considerations for Alternate or Third-Party Antenna Systems 2-35

Benefits 2-36

Limitations 2-36

CHAPTER 3**Common Wireless Medical Devices 3-1**

Overview 3-1

Wireless QoS Considerations 3-3

The Use of Unlicensed 802.11 Frequencies in Healthcare 3-5

RF Interference Common to Healthcare environments 3-6

Effects of Radar in 5GHz Band 3-10

Effects of RF on Medical Devices 3-11

HIPAA Considerations 3-12

Wireless Medical Devices 3-12

EMR Systems 3-12

Smart Infusion Pumps 3-14

Patient Monitoring Systems 3-18

Portable Radiology and Cardiology Devices 3-20

Medical Carts—COWs 3-22

RFID Devices 3-23

Communication Devices—7921/Vocera/Future of Dual Mode 3-23
 Bar Code Scanners 3-27
 Cisco Assurewave Program 3-28

CHAPTER 4

Symmetric Mobility Tunneling Deployment Guide 4-1
 Introduction 4-1
 Background on Mobility in Cisco’s Unified Wireless Network 4-1
 Symmetric Mobility Tunneling Operation 4-5
 Mobility Configuration 4-6
 Regular, Asymmetric Mobility Configuration in WCS 4-6
 Regular, Asymmetric Mobility Configuration in the Controller GUI 4-9
 Verifying Regular, Asymmetric Mobility Operation 4-11
 Symmetric Mobility Configuration in WCS 4-12
 Symmetric Mobility Configuration in the Controller GUI 4-13

APPENDIX A

Antenna Recommendations A-1
 Antenna Types and Placement Recommendations A-1

APPENDIX B

References B-1



CHAPTER 1

Wireless LAN Networks in Hospital Environments

Healthcare organizations are highly mobile by nature and hence are constantly challenged with delivering adequate wireless services to access patient records at the bedside and enhance maximum responsiveness. Many hospitals are deploying Wireless LAN (WLAN) networks to support a variety of healthcare applications. These WLANs must perform optimally in a wide variety of physically different building structures. For example, many hospitals have unconventional building shapes, diverse building materials, and connect to other buildings. WLANs also need to perform optimally with a wide variety of client devices, such as handheld devices, laptops, mobile carts, Wi-Fi voice, Wi-Fi location, etc. WLANs in hospitals must also be resilient enough to handle various sources and levels of interference.

Facing these issues, it is a fortunate site that can deploy a WLAN throughout the site with one network design philosophy. The Cisco Unified Wireless Network has the flexibility and sophistication to handle all these challenges. The Cisco Unified Wireless Network comprises the Cisco Aironet access points with antenna accessories, Cisco Wireless LAN Controllers, and the Wireless Control system. This combination of products ensures a solution that can provide optimal application performance in the most challenging RF environments.

Successfully deploying a WLAN network that is able to provide the diverse services in a challenging environment can be a challenging project for any healthcare IT organization. Doing it right the first time requires a special set of skills and knowledge that sometimes is difficult to find. By partnering with a network integrator that is Wireless Cisco certified, organizations can be confident in a deployment that meets the demanding challenges found only in healthcare environments.

Cisco's Advanced Services group can provide documented wireless installation, complete RF site surveys, and related network design and analysis services. The Advanced Services group provides a broad portfolio of services that address all aspects of deploying, operating, securing, and optimizing your network to help increase business value and return on investment.

This document is organized into the following chapters and appendixes:

- [Chapter 2, “RF Design Considerations”](#)—Discusses how a variety of factors influence WLAN network design in healthcare environments, including client type and application requirements, cell size and spectrum selection criteria, etc.
- [Chapter 3, “Common Wireless Medical Devices”](#)—Provides an overview of the wireless 802.11-based devices used in medical facilities. Both wired and wireless devices are used throughout the healthcare enterprise, therefore aggregate traffic load and network metrics must be considered in the overall wireless network design.
- [Chapter 4, “Symmetric Mobility Tunneling Deployment Guide”](#)—Describes the Symmetric Mobility Tunneling feature, including its background and deployment procedures.

- [Appendix A, “Antenna Recommendations”](#)—Describes various antenna types and their placement for healthcare environments.
- [Appendix B, “References”](#)—Contains URLs for the various documents referred to in this document.



CHAPTER 2

RF Design Considerations

Overview

The purpose of WLANs is to provide mobility for current and future applications. Often applications can be moved from the wired media to WLANs with little effort and few problems. But often the applications that were developed on the wired media do not have the connection logic or sufficiently flexible timing values to move gracefully to a WLAN, which becomes a WLAN design issue since applications have different bandwidth requirements.

The 802.11 standards are 10 years old and differ greatly from the latest 802.11 specifications. The latest 802.11n specification has protocol and hardware differences that are far different than the 1997 specification. Because the protocols and hardware of 10 years ago must interoperate with protocols and hardware simultaneously on the same RF media (air), the interoperation of the old and new becomes a design consideration. With wired networks those differences are separated and segregated by an Ethernet cable to the point of the switch port. Once the signals are in the switch, all signals become equal. Therefore, the WLAN design must consider the performance of each client device type and the application(s) that run on the client.

The Wi-Fi Alliance has a certification process for clients and access points (see the link in [Appendix B, “References”](#)). The Wi-Fi Alliance is an organization that is supported by the majority of large companies selling Wi-Fi products. The Wi-Fi Alliance’s certification processes are optional and there is no requirement for a client device or an access point to obtain Wi-Fi certifications. The vendors of Wi-Fi products pay fees to the Wi-Fi Alliance to have their devices tested for interoperability. The devices are tested to adherence to 802.11 specifications, but only to the level and to the parts of any specification that the Wi-Fi Alliance deems necessary.

It is important to consider the fact that a product that carries a Wi-Fi certification may not be of sufficient product quality to be considered enterprise class. The devices are not tested for radio signal strength and, to date, there is no certification testing for secure fast roaming, although there is an 802.11 Task Group (802.11r) working on a secure fast roaming specification. The 802.11r specification takes into account the 802.11e QoS and traffic stream specifications. Excluding 802.11r, almost all roaming decisions are made by the client and not the access point. The access points and the wireless controller have little if anything to do with client roaming. The quality of roams depends in a large part to how well a client’s firmware, radio, and the operating system work together. A client with a good integration of these attributes will have better roam times and more consistent roams.

The Wi-Fi Alliance’s Web site maintains a searchable, up-to-date list of devices that have passed their certification tests. Part of the process of evaluating the performance of clients to be added to your WLAN should include checking the Wi-Fi Web site to determine if the client is 802.11 certified and if it supports EAP and other items that may be required. The site provides information on what each certification involves. The CTIA Certification Program for cellular phones now requires that dual-mode WLAN phones go through the Wi-Fi certification for dual-mode phones.

There are applications or clients that may require RF designs that are substantially different from typical data applications. The applications that run on thin clients often need to have better RF coverage and throughput than applications designed to work in WLANs. Thin clients traditionally have not been resilient to RF retries and delays. Location applications typically need AP placement in corners of buildings. There must be sufficient interior floor access point placement to allow for accurate multi-lateration for location applications. The accuracy of the location application may also require active access points or monitor mode access points to be placed so as to form a perimeter around the floor. See [Appendix B, “References”](#) for the link to the document Wi-Fi Location Based Services Design Guide 4.1, which contains in-depth design information on location services.

The Cisco Unified Wireless Control System has a feature that imports floor plans and aids in the process of doing RF design. This planning tool helps determine AP placement based on coverage modeling. It also provides parameters that can help in coverage modeling by client performance and application requirements. To use this tool effectively, the client(s) devices' RF characteristics need to be well understood.

This chapter stresses that the coverage design of a site must consider the performance characteristics of the RF side of the client devices. You should also consider the following about the clients to ensure a good coverage design:

- The 802.11 specification version support of the client devices
- The Cisco CCX version support of the client devices
- The application requirements that run on the clients devices
- The spectrum capabilities of the clients, i.e., are they 2.4GHz, 5GHz, or both?
- The maximum data rates and transmit powers and do the clients perform dynamic transmit power control?

Another function that needs to be added to the inventory of client capabilities is what are the security capabilities. First generation clients may only support WEP encryption, while the newest devices likely support AES encryption. This is not necessarily a coverage design issue, but is a part of the overall supporting design.

Once client data is collected, more intelligent decisions can be made on the access point models, type of antennas, access point placement, and the spectrums that should be included in the design. After applications requirements are collected, more intelligent decisions can be made on the throughput requirements of the coverage area of individual access points (WLAN radio frequency cell).

The throughput of a 2.4GHz WLAN RF cell is going to be influenced by these items and others, such as:

- Beacons at data rates of 1 Mbps to support first-generation clients
- High retry rates because the cell is larger than what the client easily supports from a transmit power perspective
- Too many clients in the cell

In most cases today's 2.4GHz client radios support data rates from 1Mbps to 54Mbps. The 1 Mbps data rate provides the longest coverage distance, but in many hospital locations long distance coverage in the 2.4GHz spectrum is not usually an advantage and may in fact create poor cell throughput. The original 802.11 specification supported WLAN radio data rates of 1Mbps and 2Mbps. The modulation type (CCK) used for those data rates provides the largest coverage distance of any modulation type in an 802.11 specification to date. The highest data rates in the 802.11 specifications use the modulation type known as OFDM. The highest data rates have the smallest coverage distance. These two facts then become instrumental in cell design. High data rates have the smallest coverage area, but provide the highest throughput cells. Low data rates provide the largest coverage area, but are the lowest throughput

cells. Another disadvantage of larger 2.4GHz cells is the increased size of the RF collision domain and lower signal to noise ratios (SNR). The larger the cell, the more RF level interaction there is between WLAN clients, but also Bluetooth, microwaves, and other RF interferers.

Cell size design is more important in hospitals than any other type of site. Even newly-constructed hospitals tend to have special purpose floor plans. Floor plans vary from wing to wing and from floor to floor. Some areas have shielded walls, some areas have small private patient rooms, and some areas may be multi-patient with only curtains. Many areas have equipment that reflects WLAN signals, while others absorb WLAN signals. Some areas have a dense population of users, WLAN devices, and a high level of channel utilization. Besides the already mentioned client type and application requirements, these are considerations that factor into decisions about cell size and spectrum selection.

First determine the applications that are going to be used in all the various locations, then the throughput requirements for the locations, then the spectrums required in each location. Only then can you begin planning the access points and antenna types and their placement.

In a hospital environment, a mobile caregiver typically updates their worklist or makes calls to colleagues or ancillary departments while in motion. The caregiver expects the WLAN device to function as they move between areas. The expected service level assumes all areas of use offer continuous application availability and active voice calls. This includes elevators, which are often overlooked in RF deployments that are initiated by department-specific roll outs of wireless-enabled clinical applications. Wireless coverage in elevators can be tricky and often requires techniques that help provide coverage which can vary from building to building, depending on the elevator and building construction.

At first glance, an RF engineer may want to install an access point in the ceiling of an elevator shaft, with a directional antenna pointing down at the elevator car. In the United States, however, placement of an access point in the elevator or elevator shaft is not legal. In addition, access to hoistway shafts can only be made by authorized, certified persons. These people are typically not the same individuals who install or survey the wireless network.

The ASME A17.1, Safety Code for Elevators & Escalators section 2.8.1 states:

“Only machinery and equipment used directly in connection with the elevator shall be permitted in elevator hoistways, machinery spaces, machine rooms, control spaces, and control rooms.”

Furthermore sections 2.8.2.1 states that, “Installation of electrical equipment and wiring shall conform to NFPA 70 or CSA-C22.1.” Section 2.8.2.2 states, “Only such electrical wiring, raceways, cables, coaxial wiring, and antennas used directly in connection with the elevator, including wiring for signals, for communication with the car, for lighting, heating, air conditioning, and ventilating the car, for fire detecting systems, for pit sump pumps, and for heating and lighting the hoistway and/or the machinery space, machine room, control space, or control room shall be permitted to be installed inside the hoistway, machinery space, machine room, control space, or control room.”

The ASME 17.1 code clearly states that devices which do not comply with the regulations cannot be placed either in the elevator or hoistway (elevator shaft). As such, the RF engineer must ensure coverage in these spaces. Coverage is not only a matter of signal strength and quality, but roaming across subnets and possibly Wireless LAN Controllers, and is common as the car moves vertically between floors.

It is not possible to characterize one standard foolproof method for ensuring coverage in elevators, as there are many variables to consider. As such, a series of recommendations and best practices need to be implemented until an approach is found that provides a satisfactory level of service for each unique environment and set of wireless clients.

One common technique to provide wireless coverage in elevators is to place a diversity omni directly outside of the doors of the elevators just below the ceiling tiles. For hospitals that are deploying new 2.4GHz and 5GHz WLANs, the Cisco AIR-API131AG is a popular choice because of its attractive design and integrated diversity antenna, which radiates the signal in a downward direction. The AIR-API242AG is also another good choice when specific external antennas are required or when

access point enclosures are required. When external antennas are used, the AIR-ANT5959 2.4GHz and AIR-ANT5145-R for 5GHz are recommended when mounted below the ceiling tiles. These are colored matched to ceiling tiles and provide low-gain diversity omni-directional radiators.

In many cases, but dependent on elevator design, a closed elevator door reduces the signal inside the car by 7 to 10 dB or more. This requires that the access point or antenna be just a few feet in front of the elevator doors. To provide fast secure roaming between floors, it is best to have access points that service elevators on the same controller and to ensure that they are part of the same access point group.

The design of the WLANs mobility groups and access point groups is essential to hospital wireless deployments. See [Appendix B, “References”](#) for links to the Cisco Wireless LAN Controller Configuration Guide 4.2 and the Enterprise Mobility 4.1 Design Guide. Many hospitals will have a campus design that has Layer 3 mobility design. For VoWLAN installations configured with access point groups, symmetric mobility tunneling is crucial to prevent a dropped call as the wireless phone rapidly associates between access points on different floors. See [Chapter 3, “Common Wireless Medical Devices”](#) and [Chapter 4, “Symmetric Mobility Tunneling Deployment Guide”](#) for more information on configuration of Symmetric Mobility, which is not enabled by default. It is recommended that it be enabled for Layer 3 mobility.

When the elevator is in motion, however, it is not possible to predict how the wireless client will react to rapid crossing of cells or if it will remain in a cell. Remember that roaming is the responsibility of the WLAN clients and not the access points or supporting infrastructure. As such, wireless client roaming behavior is strongly influenced by the roaming algorithm implemented in the wireless driver supplied by the vendor. It is therefore not possible to guarantee stable connectivity in elevators due to the unique and differing environmental and wireless client characteristics.

Hospitals that are planning wireless installations today are using the AIR-AP1250, which provides backward compatibility to 802.11b/g/a wireless clients while at the same time providing 802.11n and MIMO. The new MIMO antenna technology, combined with the newer radios, has better performance in areas with high multipath as compared to non-MIMO technologies. Many areas of a hospital are prone to high levels of multipath interference due to construction techniques as well as RF shielding in some areas of the building. Recommendations for MIMO-based ceiling mount antennas are the AIR-ANT2430V-R for 2.4GHz and AIR-ANT5140V-R for 5GHz.

While this document is focused on the radio cell design, infrastructure design is equally important to a successful deployment. The Cisco Wireless LAN Controller Configuration Guide and Cisco Wireless Control System Configuration Guide provide detailed information for configuring all models of controllers and the Wireless Control Server (WCS) which manages the controllers. The Voice over Wireless LAN 4.1 Design Guide provides detailed information on the design architecture for successful VoWLAN. The Enterprise Mobility 4.1 Design Guide provides detailed information on design considerations for security, accessibility, QoS, and high availability in a single building or campus network. These documents address very important aspects of infrastructure design that create the foundation of a viable WLAN and can be downloaded at the links in [Appendix B, “References”](#)

Capacity Planning for High Bandwidth Areas

Capacity planning involves determining what areas of the hospital site use the largest number of WLAN devices and if those areas have a high uptime use. In some deployment planning models, if a device is in an area, it is considered in use and bandwidth is calculated assuming all devices and applications are running simultaneously. Also, because most WLAN devices in a hospital are wireless for mobility, plan for clients to also roam into cells and therefore a percentage of bandwidth must be reserved for roaming clients. If there are known future application deployments that may include WLAN clients, then they too should be considered in capacity planning. The 5GHz spectrum provides the most flexible method of high bandwidth design as there are 21 channels available in the 5 GHz spectrum. These channels do not have frequency overlap as is the case with the 2.4GHz spectrum. Therefore 21 access points could share

the same floor space and not interfere with each other or the clients associated to them when in 5GHz 802.11a mode. The 5GHz 802.11n wide channel mode (40Mhz) provides for 9 non-overlapping channels. When standard 20MHz wide channels in the 5GHz 802.11n band are used, there are 21 non-overlapping channels just like 802.11a.

If the high bandwidth area needs high throughput in the 2.4GHz spectrum, then the design becomes more challenging. First of all there is a high level of noise in the 2.4 GHz spectrum. Often called the noise floor, the noise is made up of the energy of all radios that are operating in the area that share the 2.4GHz frequencies that the access points are on. This energy (noise) may also come from a microwave oven, a surveillance camera on or inside the building, or a Bluetooth wireless mouse. Other sources of the energy may be WLAN devices in other areas of the hospital. 802.11 packets and Bluetooth packets that have degraded to the point that they no longer have a decodable format are considered to be noise. Any energy in the WLAN spectrum that is not decodable is considered noise. If the energy level is high enough that the nearby client or access point cannot decode WLAN packets, then it is considered interference. This is different from co-channel interference because co-channel interference involves decodable WLAN packets. Co-channel interference is really contention between overlapping cells of access points. Removing the 1Mbps and 2Mbps data rates is one technique that can greatly reduce the effects of co-channel interference.

The 2.4GHz spectrum is used by four 802.11 specifications and four 802.11 modulation types. The specifications are 802.11, 802.11b, 802.11g, and 802.11n. The 802.11 specification supports frequency hopping at 1 and 2Mbps and direct sequence at 1 and 2Mbps. The 802.11b specification supports the direct sequence data rates of 5.5 and 11Mbps. The 802.11g specification supports rates from 6 to 54Mbps and 802.11n extends those rates to 300Mbps. The 802.11g and 802.11n 2.4GHz client interoperate with each other seamlessly because they both use OFDM. However, when either 802.11 or 802.11b clients share the cell with 802.11g/n clients, then a protection mechanism is used to protect the OFDM clients from the 802.11 or 802.11b clients. This mechanism has the by-product of lowering the cell throughput. The data rates are maintained by the clients and access points, but the cell throughput is lower because of the overhead related to the protection mechanism. The protection mechanism is a 802.11 CTS packet sent at a data rate that the 802.11 and 802.11b clients can hear.

The design considerations for high bandwidth on 2.4GHz need to begin by determining the data rate requirements of the client devices. If there are clients that require 1 and 2Mbps, then getting a high throughput cell is going to be difficult to obtain. Such a cell may not get more than 7Mbps throughput if there are 802.11 clients within the cell. If there are no requirements to have data rates of 1 or 2Mbps supported, then disable them. Disabling those data rates improves the throughput of this cell and reduces the likelihood of co-channel interference in a nearby cell and contribution to the noise floor in more distant cells. A cell that has 802.11b and 802.11g support requirements may have a maximum cell throughput of 13Mbps. A cell that is 802.11g only may have a maximum cell throughput of 25Mbps. Planning above that in the 2.4GHz spectrum is far off in the future because sites are still supporting devices that are five years or older. In [Appendix A, "Antenna Recommendations,"](#) [Figure A-4](#) shows a hospital site that reduced its channel utilization from over 35% to 5% by disabling 1 2Mbps.

Another design method that is used to increase the amount of 2.4GHz throughput in a high bandwidth area is to reduce the coverage of the access point. The smaller the cell created by the access points, the more cells that can be put into an area. The most effective way to do that is reduce the transmit power of the access points. Another design consideration is to select a particular access point model or select a particular antenna type. For example, the AP1130 series has built-in antennas. By reducing the transmit power level, the cell size is reduced. The AP1130 by design has a downward signal propagation pattern with little upward signal propagation. This reduces lower floor to upper floor signal leakage, resulting in a small cell and reduced co-channel interference. If access points with cabled antenna connections are deployed, then the type of antenna selected becomes important. The typical rubber duck dipole antenna has little or no signal gain. Its signal propagation is similar to a ball and the signal is pretty much equal in all directions. These antennas result in some signal being directed to the floor above and below. With

the typical rubber duck dipole antenna, the signal level is reduced by lowering the transmit power of the access point. Antennas that clip onto the ceiling typically have a signal propagation pattern that is more downward than upward and are a good alternative to rubber ducks.

Cisco's WCS and WLC can be configured and tuned to help design such high-bandwidth areas. They also can be configured to help manage, optimize, and maintain a design that was created by more conventional methods. Links to the WCS and WLC design and configuration guides can be found in [Appendix B, "References"](#)

The 802.11n specification raises expectations of improved cell throughput, capacity, and coverage, which are all true. But what is also true is that an 802.11b client is still an 802.11b client, even when a client is associated to an 802.11n-enabled access point. The 802.11b client still has a maximum data rate of 11Mbps per second and a maximum throughput of 7.1Mbps. Its coverage may be improved by the access point's increased receiver sensitivity and increased transmit power, but the improvement is likely to be no more than 10%. In noisy environments, especially those with multipath, the throughput should be better because of the MIMO antenna technology on the access point. The result is fewer retransmissions and retries, thus increasing the throughput in the 802.11n cell. The Wi-Fi Alliance has taken the position that it will not do a certification process for 802.11n in the 2.4GHz spectrum. The 2.4GHz is a dirty spectrum with limited spectrum space. It is recommended that sites that have 802.11, 802.11b, and 802.11g clients only use 802.11n with 20Mhz channels. 802.11n has the same CTS protection mechanism that 802.11g used for mixed client cells. In a mixed cell, an 802.11n client uses 802.11n data rates, but the throughput is slowed by the protection mechanism.

The 802.11n performance is best suited for the 5GHz spectrum. There are 9 non-overlapping 40 MHz channels in 5GHz. There is only 1 in the 2.4GHz spectrum. There is no CTS protection mechanism between 802.11a and 802.11n. They both use the same OFDM modulation header. The 802.11a client can pass data in the 40MHz channel, although the 802.11a only uses half of the channel. The 802.11a client benefits by about 10% from the MIMO technology of the 802.11n access point.

In summary, the 5GHz spectrum is better suited in high-bandwidth areas than the 2.4GHz spectrum. More channels means less design work, less interference, and no non-OFDM client concerns. There is one planning issue relating to clients. It must be determined what 5GHz channels the clients support. Many 802.11a clients only support 12 of the current 21 channels. If you have clients that do not support all 21 channels, then those channels must be removed from the controller configuration.

Channel Planning for High Density Areas

The first item to determine in coverage areas with a high density of users is how many of them are WLAN users. Then determine the applications they use and the required bandwidth. The rule of thumb for the number of users per access point has been 15 to 25, which was originally associated with 802.11b. The current standards have higher data rates and the current access points have both 2.4GHz and 5GHz radios with a 100Mb Ethernet for the AP1130 and AP1240 series. The AP1250 series has a 1Gig Ethernet port. However that rule of thumb is still fairly accurate. Current clients are typically 802.11g and 802.11a with data rate support of 54Mbps. There will be enterprise level 802.11n clients with 300Mbps data rate support in the near term. The advantage is that the current access point can already simultaneously support the 2.4GHz and 5GHz clients, which means the access point has the processing power and buffer space to support 15 to 25 users on each access point radio.

Channel planning in high density areas involves more than simply determining the number of users that can use the access point(s). For VoWLAN 802.11b clients that do not use the latest 802.11e specifications, channel planning would be based on the old seven active G.711 streams or seven VoWLAN clients per access point. If the area still has old data client technologies that are not 802.11e aware or are not doing WLAN QoS, then the number of users per access point and channel must be fewer than if all the clients in the coverage area are 802.11e aware. Currently there are only a few data clients and VoWLAN clients that are WMM QoS aware from the client to the access point.

To design a good channel plan with little or no client support for 802.11e QoS or the Wi-Fi Alliance WMM, the high density area requires even more traffic analysis and capacity planning. The traffic analysis would include the data rates supported by the clients. For data clients, the higher the data rate supported by the client radio, the less time a packet is on the channel and therefore the lower the channel utilization. For VoWLAN, the higher the data rate supported by the voice clients, the more call legs that can be recreated on the channel with less negative impact on the quality of calls on the channel. It needs to be noted that voice packets are small, which is by design for VoIP. When small voice packets are dropped or lost, the call quality is better than if large packets are dropped or lost. The 802.11 protocol overhead and the small size of VoIP packets affects a channel's highest throughput rate. A VoWLAN client using the G.711 codec has the same throughput at the 36Mbps data rate as the 54Mbps data rate.

Channel planning should not use G.711 call streams per access point, but rather call streams per radio channel. The typical number of calls achievable with good design and client behavior in channels with good SNR values is:

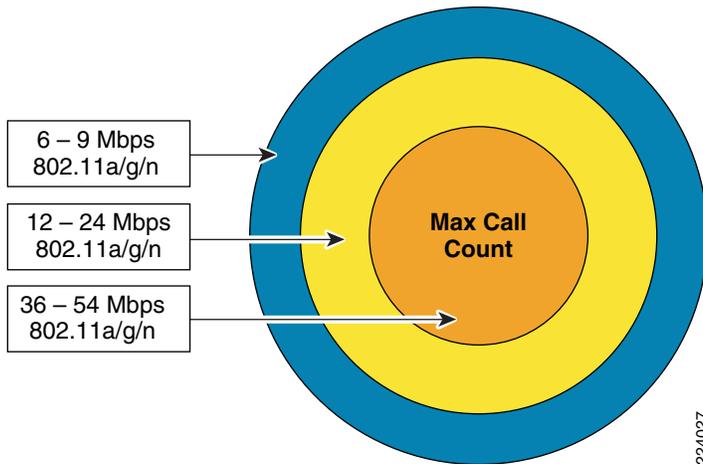
- 7 active streams with 802.11b only
- 10 active streams with 802.11b/g
- 14 active streams with 802.11g only
- 20 active streams 802.11a only
- 22 active streams 802.11a/n

VoWLAN planning for high density should be done with clients that support call admission control (CAC) and QoS. CAC protects both the radio channel from over subscription and the call quality of active calls. When the utilization of the radio channel does not support an additional call without negatively affecting active calls, then additional calls are prevented. This is a very important feature supported by the Cisco Unified Wireless Network.

Figure 2-1 shows the data rates at which VoWLAN achieves the highest number of calls relative to cell size. When an access point is configured with low data rates (i.e., 6Mbps and 9Mbps) it effectively has a larger cell coverage area than an access point that has those rates disabled. The larger cell size allows more VoWLAN and data clients to be active with an access point. The cell likely has a higher channel utilization and noise floor and therefore supports fewer VoWLAN calls. A cell that supports data rates of 6Mbps through 54Mbps supports fewer calls than a cell that has only the 36Mbps to 54Mbps enabled. The packets of the clients in the 6Mbps coverage area of the cell have longer air time than the packets sent at data rates of 12 – 24Mbps and longer than those packets sent at 36 – 54Mbps. Figure 2-1 shows a cell with the combination of 54Mbps client and 6Mbps clients. The 54Mbps clients are slowed slightly by holding off for the 6Mbps clients.

This then becomes part of the criteria for call planning. The original Cisco guidelines for the call planning stated 7 calls per access point, however that design logic is no longer valid. With the advent of 802.11e and WMM, the design criteria is call streams per RF channel. A call stream is, for example, a Cisco 7921G VoWLAN phone calling a wired desk phone. Two call streams would be two 7921Gs calling two wired desk phones or two 7921Gs calling each other. If the two 7921Gs that are calling each other are associated to the same access point, then those two call streams are in the same cell and are on the same RF channel. The new call planning criteria is based on the number of call streams on the same RF channel. It is important to remember that RF channels can overlap. Two access points in near proximity of each other can share the same RF channel and therefore the number of call streams suitable for the channel, based on call admission control logic.

Figure 2-1 Data Rates and Highest Number of Calls Relative to Cell Size



The recommended configuration for design for VoWLAN depends on the client's level of support for the QoS basic service set (QBSS) or WMM. The configurations for these options are explained in detail in the WLC design guide.

It is recommended that all SSIDs be configured using WMM. [Figure 2-2](#) shows the four access categories defined as 802.11e and used by WMM. Using the categories for the applications as defined in the description field ensures that voice packets get QoS priorities over data transactions.

If the SSID is not configured for WMM and QoS, then the traffic for these clients uses standard DCF frames, which is equivalent to Best Effort.

Figure 2-2 Defined Access Categories

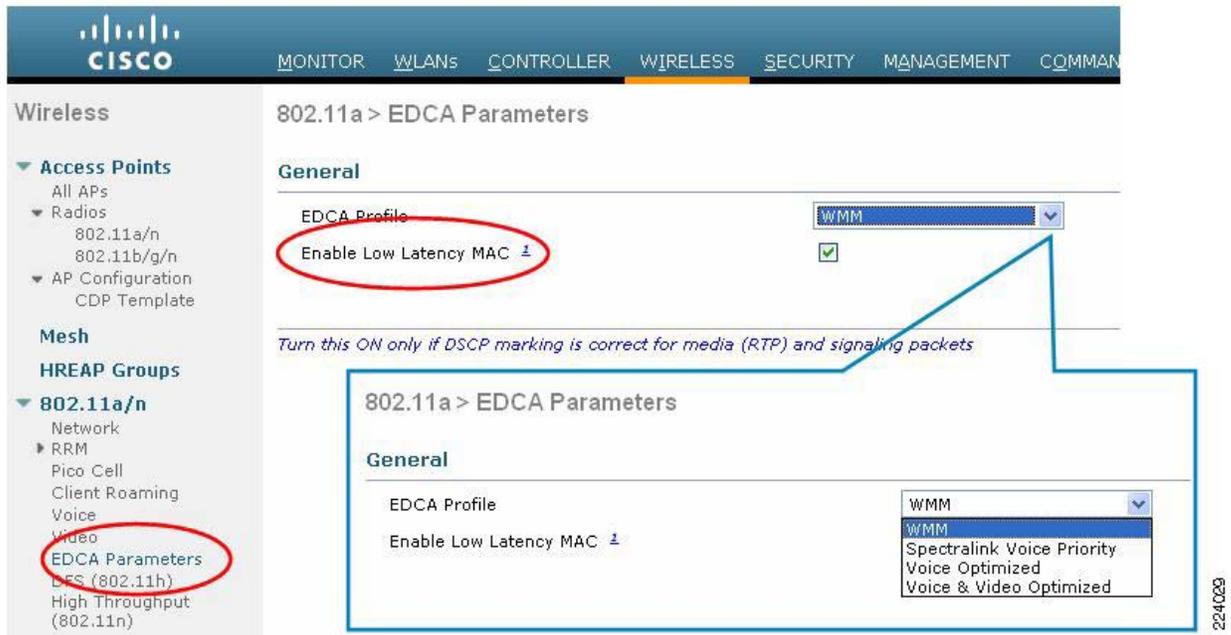
The screenshot shows the Cisco Wireless configuration interface. The top navigation bar includes MONITOR, WLANs, CONTROLLER, WIRELESS (highlighted), and SECURITY. The left sidebar is titled 'Wireless' and contains a tree view with categories like Access Points, Mesh, HREAP Groups, 802.11a/n, 802.11b/g/n, Country, Timers, and QoS. The 'QoS' category is expanded, and 'Profiles' is circled in red. The main content area is titled 'QoS Profiles' and displays a table with the following data:

Profile Name	Description
bronze	For Background
gold	For Video Applications
platinum	For Voice Applications
silver	For Best Effort

22-4028

To further facilitate high performance cell capabilities for various VoWLAN clients and video requirements, WLC provides specialized EDCA options (see Figure 2-3).

Figure 2-3 EDCA Parameters



Designing Access Point RF Coverage to Meet the Needs of Wireless Applications

Hospitals are likely to have the widest variety of mobile applications of any industry. Patient monitors, mobile ultrasound systems, barcode scanners, PDAs, location tags, VoWLAN, and guest or patient laptops may all be communicating through the same access point. Of this group of WLAN devices, the two that require the most attention to WLAN design are the VoWLAN and location or RFID tags that use the WLAN 2.4GHz spectrum. All location or RFID tags that work on a WLAN are using 2.4GHz. Location tags can be tracked by Wi-Fi and chokepoints.

A chokepoint location system consists of multiple-frequency tags (Wi-Fi 2.4-GHz and lower frequency transmitting devices) and electronic devices like Exciter and WherePorts, known as chokepoint triggers. The chokepoint triggers use 125kHz frequencies to generate magnetic fields. Chokepoint triggers are the electronic devices placed in chokepoints. Chokepoints can be placed in numerous locations in a hospital, such as a building entrance, an entrance to an operating room, or a workstation in a manufacturing facility. When the tag passes within range of the chokepoint, the low-frequency field awakens the tag, which then sends a message over the Cisco Unified Wireless Network. The 802.11 information packet includes the chokepoint device ID and can also include sensor information (such as temperature or pressure), depending on the type of tag that is transmitting. The packets can be transmitted on a per-event basis or may be configured for long transmit intervals when not mobile. Typically they use motion sensors to detect motion and begin the migration to short transmit intervals. Thus, while a tag may transmit 802.11 packets at a 10 minute rate when not in motion, when in motion that same tag may transmit every 30 seconds. If that same tag encounters a chokepoint trigger, it transmits immediately regardless of the transmit interval currently in effect. Even at 30 seconds intervals the tags are

transmitting at intervals that do not affect the throughput of a channel significantly. Active RFID tags typically transmit at +18dBm on 2.4 GHz by default, which would not be characterized as low power. But it is the long intervals that do not affect the throughput of a channel to the point that it becomes a design issue that impacts the effect the signal level has on other applications.

A chokepoint location system provides precise detection (covering areas ranging from a few inches to twenty feet, depending on the vendor). It needs to be noted that chokepoint triggers and asset tags from different asset tag manufacturers are **not** compatible with one another. Thus, it is recommended that chokepoint triggers and asset tags be purchased from the same vendor. All Wi-Fi vendor tags do use 2.4GHz to transmit to the access point. The access points hear the packets from the tags, even if the access points have the data rates used by the tags disabled. The Cisco Wireless Control System (WCS) provides location maps of clients, Wi-Fi tags, and chokepoints. The WCS also provides a planning tool for location that aids in the placement of access points. WCS allows for the reference placement of chokepoints on the location maps. Chapter 5 of the WCS guide details how to use that function. Links to the WCS guide and the Wi-Fi Location Based Service Design Guide 4.1, which contains design recommendations for location services, are in [Appendix B, “References.”](#)

It has been noted how the beacons of 1Mbps and 2Mbps adversely affect the throughput of a cell and nearby cells on the same frequency. As an option, the WLCs and WCSs can manage selected access points to be 2.4GHz monitors, which means that the 2.4GHz radios of monitor access points only receive packets and do not transmit packets. Radios in monitor mode do not transmit beacons or any other sort of packet or signal. As monitors they do not adversely affect the throughput of nearby cells. The monitor mode access point can hear the location tags and forward to the location servers the packets from the tag and therefore the information from the tags. The monitor mode access point provides the information for triangulating the location of a tag just as fully active access points would, but without adding any RF signals.

**Note**

Triangulation is a generic term used to determine the location of devices from three reception points. The actual method used to determine the location of an asset is tri-lateration when the tag is detected by three access points and multi-lateration when detected by more.

There are additional features that are provided by monitor mode, including rogue client/access points and IDS.

VoWLAN clients like Vocera require a dense deployment of 2.4GHz access points. The Vocera clients typically transmit at lower power than PC-based devices or barcode scanners. For their voice quality to be acceptable and for roaming from one area to another to have acceptable voice quality, the access point deployment needs to be carefully designed. The basic guideline for Vocera 802.11b clients is to provide a cell that has an 11Mbps data rate a cell edge power level of -67dBm. The other cell design issue is that the 802.11b cells have a 20% cell overlap with the next 802.11b cell. The WCS has a pre-configured design tool set up for the Cisco 7920 phone. The design requirements for the Vocera Badge and 7920 are the same, so pre-configured 7920 designs can be used for the Vocera 802.11b. The planning mode feature for access points to support VoWLAN clients is in chapter 5 of the Cisco Wireless Control System Configuration Guide (see [Appendix B, “References”](#) for the link).

The VoWLAN planning feature tool allows the parameter changes in the tool to match up to the performance characteristics of the VoWLAN device. The tool provides a map, based on signal propagation algorithms, of what the coverage would be for that VoWLAN client. Access points can be dropped on the map or removed, thereby providing an excellent method of estimating access point placement and access point qualities required to provide three different levels of call quality support. Again, the better that the radio performance of the client is known, the better this tool can contribute to WLAN design.

The WCS coverage planning tool can be used for data clients. By using the tool as if the data parameters were entered for a phone when they are actually the data parameters of the another client type, the planning tool can provide coverage estimates for non-phones. This tool could be used for devices or applications that have little performance or retry resiliency (see [Figure 2-4](#)).

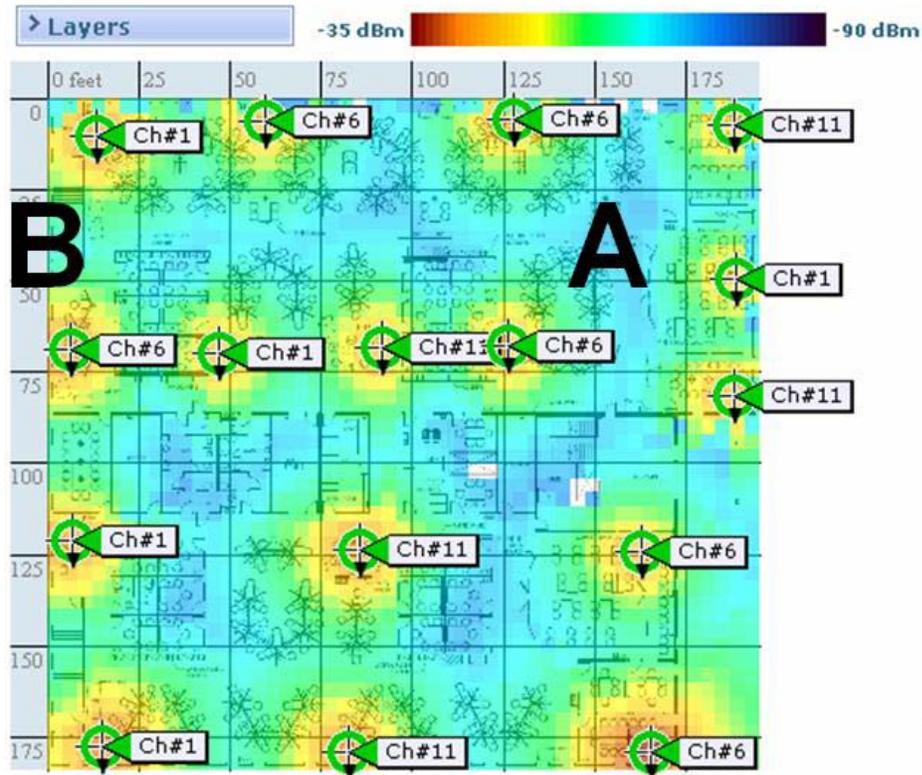
Design Considerations for the Location Services and Voice in the Same Site

VoWLAN and tag location services running in the same site or coverage areas need to be designed. The coverage needs to be designed to the throughput requirements of the VoWLAN client and the accuracy requirements of the location application. The location design guide and the location appliance deployment guide links are in [Appendix B, “References”](#) Access point placement recommendations are in chapter 1 of the guide.

Location services design differs from traditional WLAN designs in that the antenna type and access point placement are designed with the criteria of achieving a three access point triangulation of the asset tag. To have accurate reporting of the location of an asset tag, the tag needs to be seen by at least three access points that are on the coverage perimeter of the asset tag. Traditional antenna type and access point placement designs were about the cell design needed to provide the proper amount of signal for the applications.

Tag A's location is triangulated by at least five access points that are within the signal perimeter of the tag. Tag B has only two access points that are with the signal perimeter of it. And because there is no access point to the left of Tag B, the reported location of the tag may be anywhere on the upper left corner of the floor. Tag A was reported within 18 inches of its actual location. Results vary from site to site as this is just a tested example.

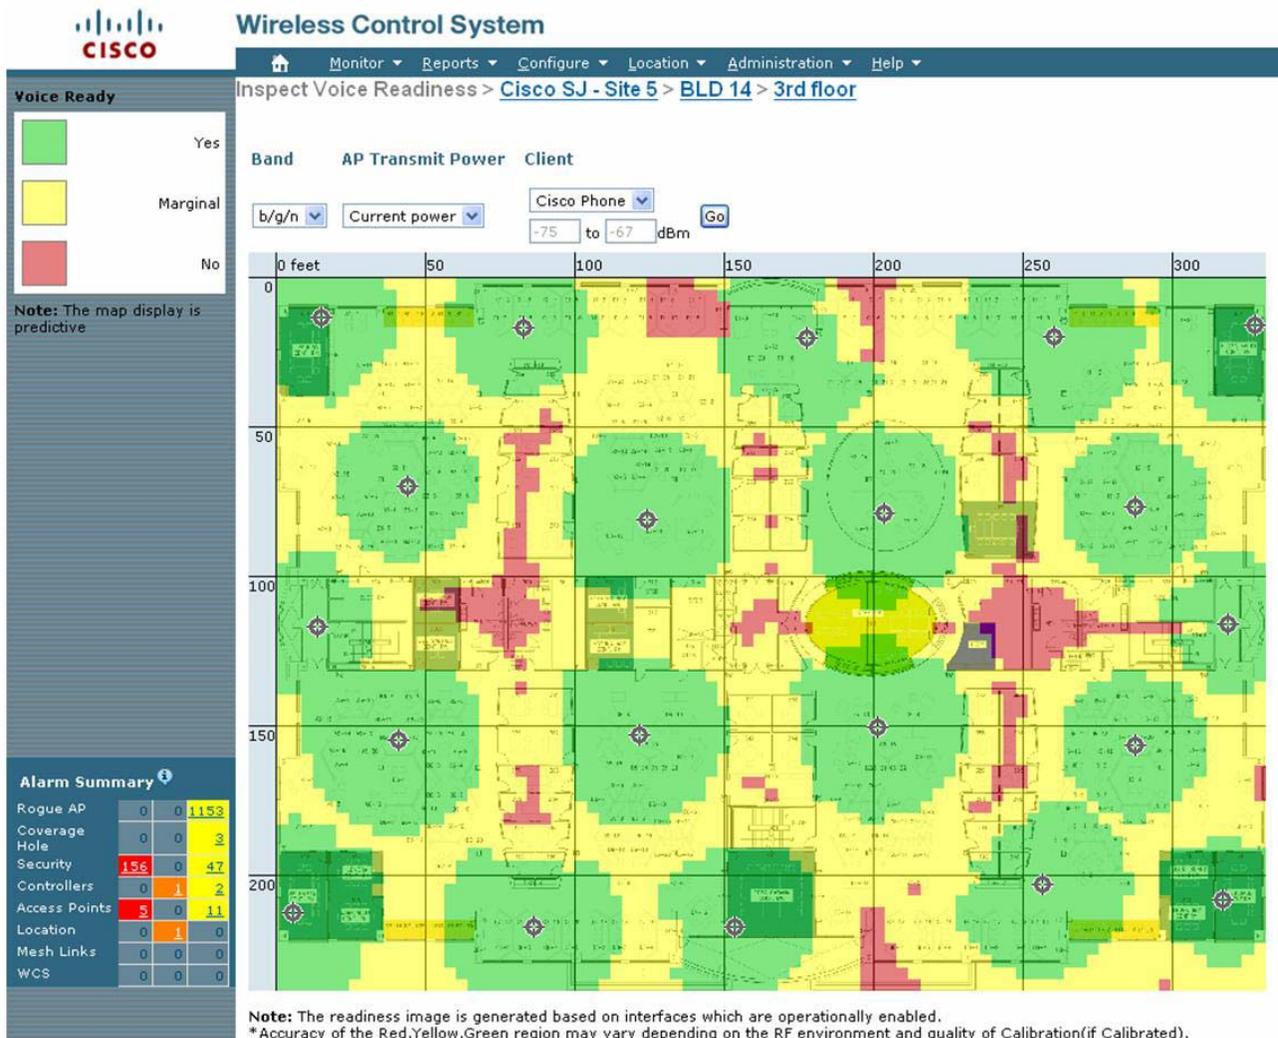
Figure 2-4 Access Point Design Example #1



22-403-1

The Figure 2-5 access point design is noticeably different than traditional deployments. The access points are placed in the corners of the floor and along the outer walls. That is following the recommendation of the location services design guide that the access point deployment be such that they form a complete floor perimeter. See Appendix B, “References” for links to location design guides. When the access points are in the corners and floor edges, it is more likely that the asset tag is between or inside the location of three access points. For the location algorithm to have a reasonable amount of accuracy, the location tag, VoWLAN phone, or patient monitor must be inside three access points.

Figure 2-5 Access Point Design Example #2



The location design guide should be consulted for the latest design recommendations for tags. The above design was per the 50-70 linear feet recommendation. In the floor space marked 100 feet to 175 feet, the APs are about 70 feet apart. The assets tags must be seen by the access points with a signal level of -75dBm. The above location design was tested with 802.11b/g and 802.11a VoWLAN clients. The VoWLAN test results did not indicate any particular behavioral problems for VoWLAN clients. Traditional data deployments and VoWLAN deployments have used a stagger design when a building had a multi-floor structure and design requirement. Typically the access points were not deployed on a stacked vertical deployment except at elevators. If access points are developed following the location guide, then at a minimum there is a vertical access point deployment in each corner.

VoWLAN deployments, like asset tag location deployments, have a denser population of access points than most data-oriented application deployments. When data deployments are mixed with VoWLAN and location, there is a radio frequency collision domain problem. To manage the collision domain, it more than likely requires the management of 802.11b/g data rates. Management of the transmit powers manually may be required as RRM is not likely to be enough. Review [Capacity Planning for High Bandwidth Areas](#) and see [Figure A-4 in Appendix A, "Antenna Recommendations."](#) RRM configurations can be found in the WLC and WCS guides.

Assets tags are not effected by disabled data rates or transmit powers. Asset tags, by design, typically transmit at a data rate of 1Mbps and with low transmit power. The 1Mbps data rate has the longest transmit range of any data rate defined in an 802.11 specification. The packets are small because they contain little data. By sending small packets at low transmit powers, the tags maintain their optimum battery life. Asset tags may not send that any more often than once every ten minutes, so they are not an influence on the throughput of the WLAN cell. Even though the Cisco access points may have the data rates of 1, 2, and 5.5 disabled, the access points still hear and process the packets sent by the asset tags. When those rates are disabled, the access points do not advertise support for those data rates in beacons and probe responses.

Access point density to support 90% accuracy for asset tags to 3 meters should be all that is required in a Vocera 11b or 7920 11b deployment. It may be possible to have access points in a monitor only mode and still support those clients in the same cells as the asset tags, if the cells sizes match the 7920 design guideline of a -67dBm cell edge for 11Mbps RTP packets. The Vocera badges and 7920 guides recommend a transmit power of 20mW on the access points.

Figure 2-6 shows the cell coverage for a sample hospital floor with WLAN data-only applications.

Figure 2-6 Cell Coverage for a Sample Hospital Floor with WLAN Data Only Applications

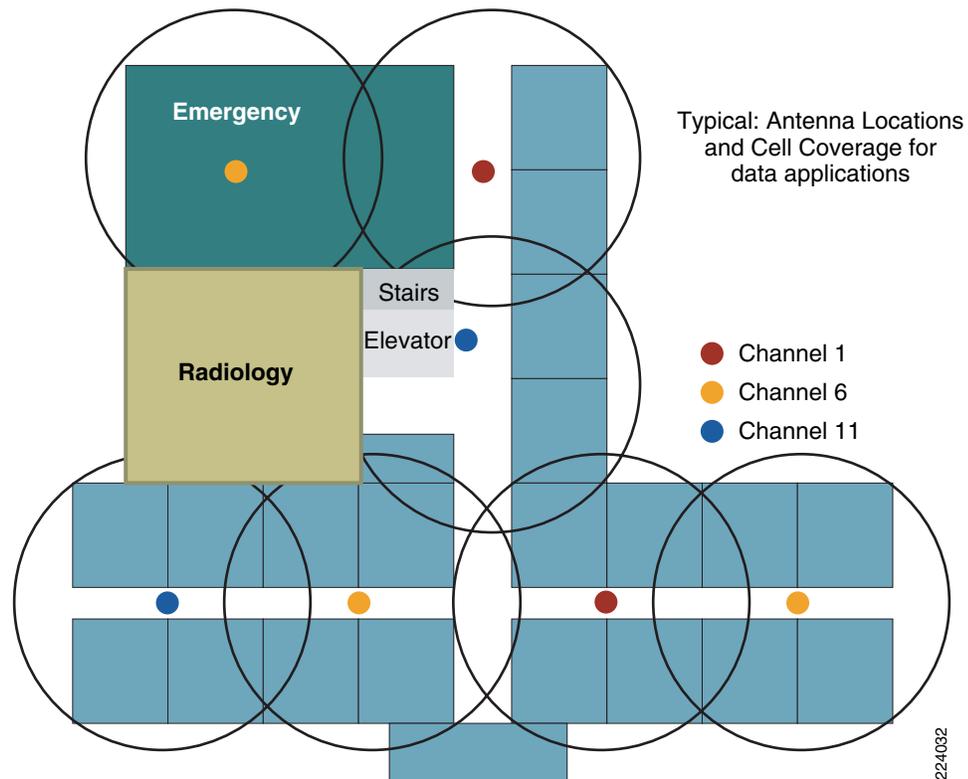
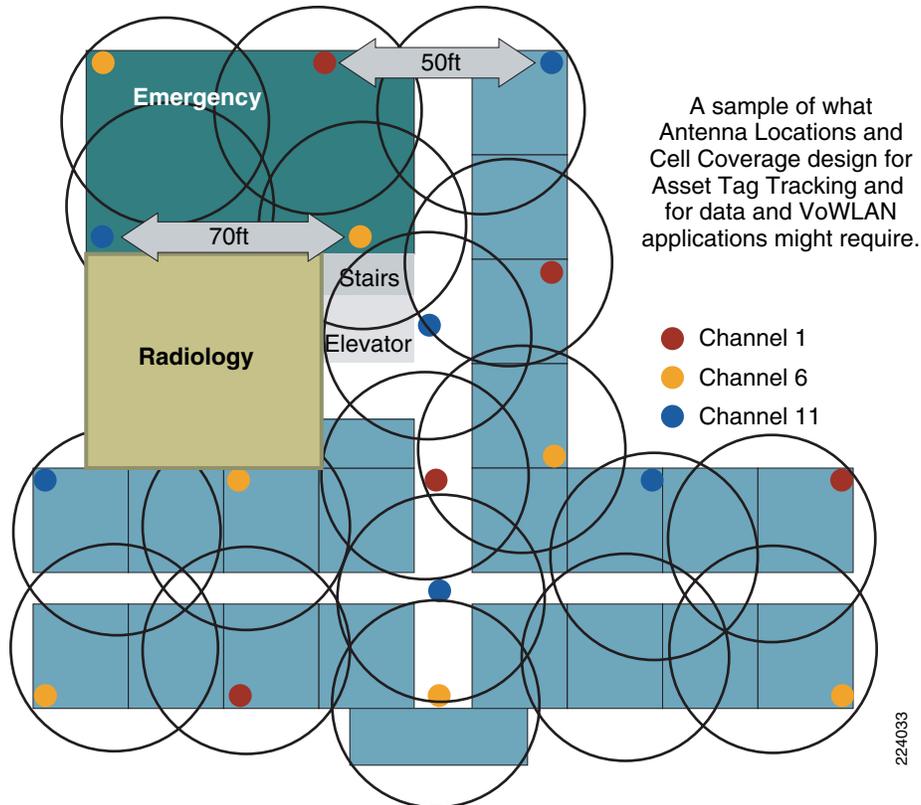


Figure 2-7 shows the cell coverage for a sample hospital floor that supports WLAN data, asset tags, and 802.11b/g VoWLAN clients. Figure 2-7 does not include chokepoint for asset tags.

Figure 2-7 Cell Coverage for a Sample Hospital Floor Supporting WLAN Data, Asset Tags, and 802.11b/g VoWLAN Clients



Design Considerations for 802.11b/g/n and 802.11a/n in the Same Coverage Area

Access point density to support 90% accuracy for asset tags to 3 meters should be all that is required to support VoWLAN clients on 802.11a/n. The access point placement that supports the asset tags at a -75dBm deployment supports VoWLAN 11a clients if those have transmit power levels of 20mW or more. In the last couple of years 802.11a deployments have become more common for two important reasons:

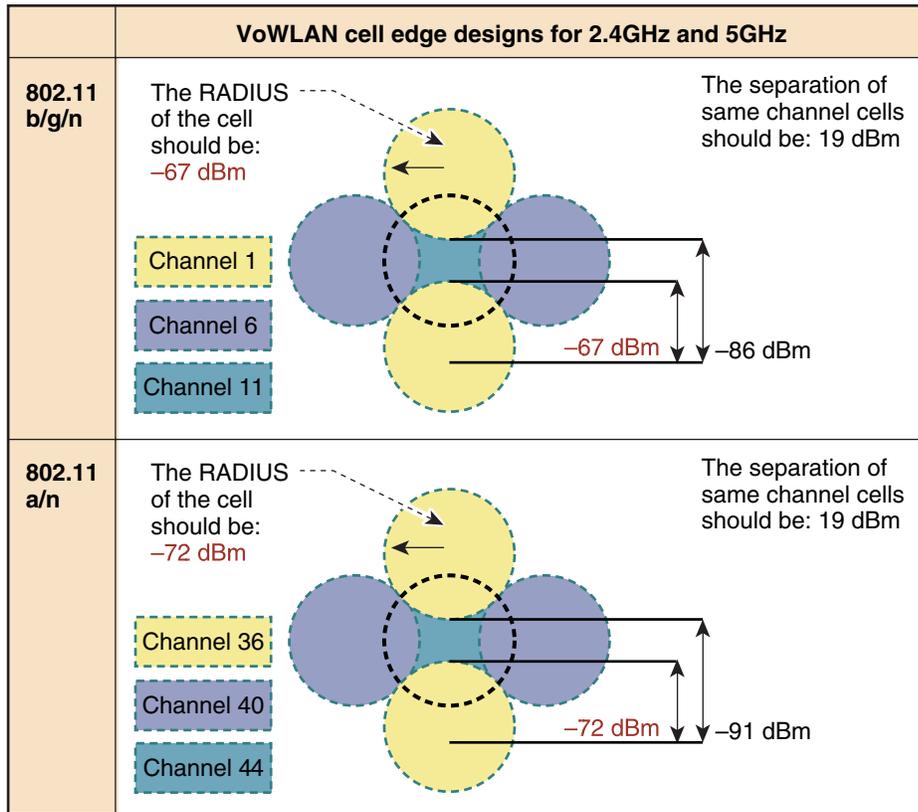
- More laptops have 802.11a radio support.
- Radios have better performance than earlier 802.11a radios.

There are more channels available now and the Cisco AP1130, AP1240, and AP1250 support the additional channels. The AP1250 supports all 21 802.11a channels while the AP1130 and AP1240 support 20 802.11a channels. The 802.11a spectrum has been avoided to some degree in the past because most clients and access points did not have transmit powers greater than 30mW. That was typically 10mW less than 802.11g and 20mW less than most 802.11b clients. The different cell performance was therefore a concern. With the high density of the today's deployments, many sites have 802.11b transmit power of 30mW or less configured. The design guidance for Vocera 11b and 7920 is 20mW. The 802.11a cell of about 30mW would be larger than that at many sites. The cell would also have, in many cases,

10dB better signal to noise ratio (SNR) on 802.11a channels than on 802.11b channels. There are also fewer devices deployed in the 5GHz spectrum that would interfere with 802.11a WLAN clients. The better the SNR, the larger the cell and even with devices that may have less transmit power. It is the ability to receive decoded packets at the radio of the access point and the client at packet error rates of 1% that determines the size of the VoWLAN cell.

With denser 2.4GHz cells at lower transmit powers and limited support for lower 802.11 data rates, it is more likely that the access point configured to support these dense 2.4GHz cells can also support a VoWLAN 5GHz cell. Therefore the RF design for asset tags may well be the design that supports the VoWLAN 5GHz cell. The 5GHz needs to be tested with the VoWLAN clients that are used in that network. The 5GHz cells need to be verified to by the -72dBm cell edge at 18Mbps (see Figure 2-8).

Figure 2-8 VoWLAN Cell Edge Designs



224034

Determining the Cell Edge of the Clients

The cell edge and same channel signal separation presented in Figure 2-8 is a design recommendation representative of ideal conditions. These ideal conditions are very difficult to achieve in real-world deployments, but should still serve as a design goal.

Each client type is going to have different radio performance capabilities. The Cisco 7920 phone has different radio performance than the Cisco 7921 phone, and they are both different than the Vocera badge. They have different antennas, transmit powers, and levels of CCX support. How well they each receive data from the access point is also an important consideration in determining each device's

coverage area. But how well the access point receives their individual signal is the most telling characteristic of their application performance. If their packets do not reach the access point, then the access point does not send application data back to the clients.

How is the performance of clients determined and what is measured? Take the different wireless clients to the same location to be tested, then take them to what may be the cell edge of a know location. Attach all of them to same access point and the same SSID. Pass active data to the devices (ping is still a good test application). The test access should be configured with only a 802.11b data rate of 11Mbps for determining the 2.4GHz cell edge. It should also be configured with only a 18Mbps data rate for the 5GHz cell edge. Then use either the GUI or CLI commands on the controller to determine the RSSI levels of each of the clients. The following is the CLI capture of a client; much of the same information is available on the GUI of the WLC.

```
(WiSM-slot4-1) show client detail 0018ba78c444
Client MAC Address..... 00:18:ba:78:c4:44
Client Username ..... unknown
AP MAC Address..... 00:15:c7:fd:b2:60
Client State..... Associated
Wireless LAN Id..... 2
BSSID..... 00:15:c7:fd:b2:61
Channel..... 6
IP Address..... 10.30.0.109
Association Id..... 5
Authentication Algorithm..... Open System
Reason Code..... 0
Status Code..... 0
Session Timeout..... 0
Client CCX version..... 4
Client E2E version..... No E2E support
Mirroring..... Disabled
QoS Level..... Platinum
Diff Serv Code Point (DSCP)..... disabled
802.1P Priority Tag..... 6
WMM Support..... Enabled
U-APSD Support..... Enabled
  U-APSD value..... 15
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
Encryption Cipher..... None
Management Frame Protection..... No
EAP Type..... Unknown
Interface..... voice
VLAN..... 30
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..... 679107937
  Number of Bytes Sent..... 715952059
  Number of Packets Received..... 3231420
  Number of Packets Sent..... 3733073
  Number of Policy Errors..... 0
  Radio Signal Strength Indicator..... -47 dBm
```

```

Signal to Noise Ratio..... 27 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]

AP.4e68(slot 0) .....
  antenna0: 5 seconds ago -50 dBm
  antenna1: 34 seconds ago -44 dBm
AP1250-0017.94cc.(slot 0) .....
  antenna0: 499760 seconds ago -53 dBm
  antenna1: 499760 seconds ago - 51 dBm

```

This shows how the client device is seen at two access points.

The following is the data collected from a access point from 4 the 802.11b/g clients in the same location at the same time:

```

Cisco 7921 - RSSI -55 SNR 51
Blackberry 8820 - RSSI -49 SNR 36
Cisco Aironet CB21 - RSSI -49 SNR 28
Intel 3945 a/b/g - RSSI -57 SNR 28

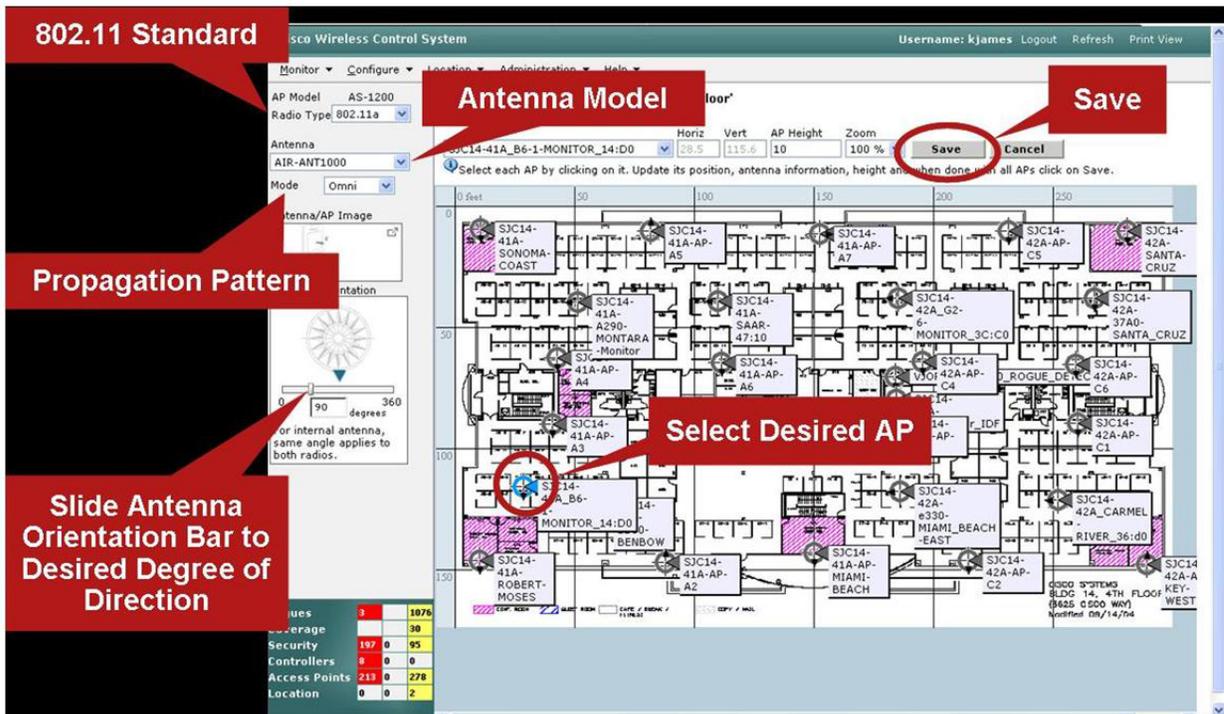
```

If these were the four clients to be used in a site, the cell design should follow the client with the lowest RSSI value. In this case the cell design should be done to the Intel 3945, but this does not mean that the Intel radio is the worst performing radio of the group. In this case that radio was imbedded in a mobile device. The complete physical design of the device influences the RSSI value seen at the access point.

Using the Wireless Control System to Estimate the Number of Access Points to Support Applications

The WCS has an access point planning feature. The planning uses floor maps imported into the systems from jpeg type formats or CAD files. The floor maps are then given floor names and building names so that the maps can be stored on the system. Once created, the floor maps can be used for the access point planning feature, readiness reporting, or for audits (see [Figure 2-9](#)).

Figure 2-9 Imported Floor Map



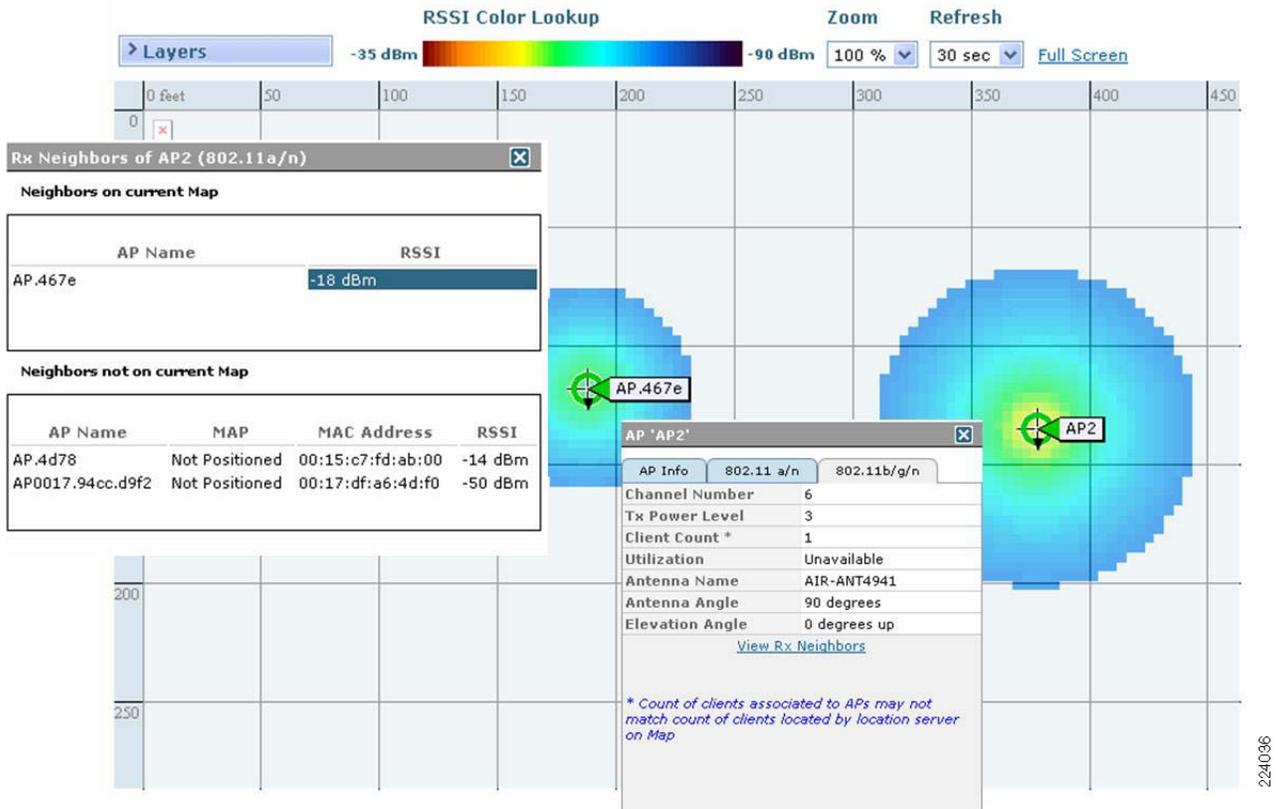
The planning tool allows for the selection of access point types, the 2.4GHz or 5GHz spectrum, and the antenna type. In a hospital, the antenna should always be a diversity antenna solution that mounts just below the ceiling tile. As per the earlier cell design discussion, the cells in most areas of a hospital are dense. Therefore when using this tool, the selected antenna should be a lower gain antenna (< 5dBi). If it is known prior to the use of this tool what antenna will be used and where it will be used, then use that information with the planning tool. Optionally, different antenna types can be used with the tool to see which antenna type may be best for the floor or an area of the floor. The tool is quite flexible. For a better understanding of the tool, see chapter 5 of the WCS configuration guide (see the link in [Appendix B, “References”](#)).

It is recommended that once the map is loaded into WCS, that the map be calibrated. That process includes accounting for walls and other building obstructions. By calibrating the map, the signal propagation algorithms more accurately predict signal coverage, VoWLAN readiness, and asset tag locations.

The planning tool includes functions for adding, positioning, and removing access points from the floor, adding, editing, and deleting floor space, and then re-computing the coverage prediction. The effectiveness improves as more is known about the radio performance of the clients, the size of the cell required for bandwidth requirements, and user density. By using the definition of service options with the above described client performance data, a more accurate map can be produced for your estimated number of access points. Optionally, the planning can be set from a safe to an aggressive number of access points.

[Figure 2-10](#) shows two access points that are added to the map. AP2 is shown as having the channel number assign of 6, a power level of 3, and one VoWLAN client associated to it. The second popup window shows that there are two other access points in the floor area. These could be access points that are valid to the network but not yet assigned or they could be rogue access points. What is shown in this planning tool is the MAC address of the two access points and their RSSI values. Both of these two access points are very near the access point known at AP.467e.

Figure 2-10 Two Access Points Added to a Map



When using this tool for location services there is also the option of adding asset tag choke points.

Using the Wireless Control System to Maintain High Quality VoWLAN Performance

A VoWLAN network in a hospital requires management and maintenance. The most automated tool to do that is the WCS. The WCS provides a variety of reports. For example, an audit report can run against the network configuration and can be scheduled to be run at a particular time and date. The information in the report can be collected on an hourly, daily, or weekly basis. The WCS manages the running of the reports. Once run, reports can be saved to a CSV or PDF file format. That file can then be forwarded to another system via ftp or E-mail.

Cisco Spectrum Expert Wi-Fi integrates with WCS for real-time spectrum intelligence for Wi-Fi and non-Wi-Fi interference sources that may influence VoWLAN performance. The reports indicate more than the presence of a microwave oven or a Bluetooth interferer; they also show the percentage of the channel utilized by the interferer. WCS reports type, discovery time, affected channels, and power levels. The spectrum expert sends information to WCS so that it can create reports, which also can be forwarded to another system via ftp or E-mail. See chapter 9 of the WCS configuration guide (the link is in Appendix B, “References”).

Figure 2-11 shows a reported microwave interferer. By clicking on the case ID Number, the MAC address and IP address of the reporting access point is displayed. The signal strength of -42dBm is quite high. The microwave has either a very bad leak or the access point is within 15 feet of the microwave.

A VoWLAN call through this access point is most likely going to drop if the call is simultaneous with the microwave running. The use of channels 6 and 11 should be avoided around microwaves. A microwave with this behavior needs to be replaced.

Figure 2-11 Reported Microwave Interferer

The screenshot shows the Cisco WCS interface with the following components:

- Navigation:** Home, Monitor, Reports, Configure, Location, Administration, Help.
- Search:** Quick Search (<IP, Name, SSID> Go), Search Alarms, New Search...
- Alarm Summary:** A table showing counts for various categories: Coverage Hole (0, 0, 3), Security (156, 0, 47), Controllers (0, 1, 1), Access Points (65, 0, 22), Location (0, 1, 1), Mesh Links (0, 0, 0), WCS (0, 0, 0).
- Interference Alarms Table:**

Detected By	ID	Device Category	Power (dBm)	Nearest AP 1	Channels Affected	Acknowledged
<input type="checkbox"/>	171.70.221.252	CSE:1197401843/2054	Microwave Ovens -56	00:17:df:a6:42:d1		No
<input type="checkbox"/>	171.70.221.252	CSE:1197401843/1796	Microwave Ovens -42	00:17:df:a6:52:10	10..12	No
- Tooltip:** Interference of category Microwave Ovens, power -42 dBm detected near 00:17:df:a6:52:10 by Spectrum Expert: 171.70.221.252 (171.70.221.252), ID = CSE:1197401843/1796

Figure 2-12 is a sample of the WCS condition reporting the helps identify conditions that may have caused VoWLAN problems. This screen is from the WCS > All Alarms > Severity Configuration screen.

Figure 2-12 Sample WCS Condition Reporting

The screenshot shows the 'Severity Configuration' screen with the following components:

- Alarm Summary:** A table showing counts for various categories: Rogue AP (0, 0, 1031), Coverage Hole (0, 0, 3), Security (156, 0, 47), Controllers (0, 1, 1), Access Points (65, 0, 30), Location (0, 1, 1), Mesh Links (0, 0, 0), WCS (0, 0, 0).
- Configurable Alarm Conditions Table:**

Condition	Category	Severity
<input type="checkbox"/> Mobility anchor control path down	Controllers	Major
<input type="checkbox"/> Mobility anchor data path down	Controllers	Major
<input type="checkbox"/> Multiple users	Controllers	Informational
<input type="checkbox"/> Network disabled	Controllers	Informational
<input type="checkbox"/> Network enabled	Controllers	Informational
<input type="checkbox"/> POE Controller failure	Controllers	Critical
<input type="checkbox"/> Power supply change	Controllers	Critical
<input type="checkbox"/> RRM grouping done	RRM	Informational
<input type="checkbox"/> Radar channel detected	Access Points	Informational
<input type="checkbox"/> Radio administratively up and operationally down	Access Points	Critical
<input type="checkbox"/> Radio core dump	Access Points	Informational
<input type="checkbox"/> Radio coverage threshold violation	Coverage Hole	Minor
<input type="checkbox"/> Radio current channel changed	Access Points	Informational
<input type="checkbox"/> Radio interface regulatory domain mismatch	Access Points	Critical
<input type="checkbox"/> Radio interference threshold violation	Access Points	Minor
<input type="checkbox"/> Radio load threshold violation	Access Points	Minor
<input type="checkbox"/> Radio noise threshold violation	Access Points	Minor
<input type="checkbox"/> Radio transmit power level changed	Access Points	Informational
<input type="checkbox"/> Radiocard failure	Access Points	Critical
<input type="checkbox"/> Radios exceeded	Controllers	Major
<input type="checkbox"/> Radius Server WLAN deactivated	Controllers	Major

The first item in this example is “Mobility anchor control path down” and it is listed as having a priority of major. In a hospital that has mobility groups created in a floor-by-floor configuration, roaming between floors is critical. Hence this item should be changed to critical. This snap shot of configurable alarms shows a dozen items that affect VoWLAN performance.

The actual VoWLAN client call performance can also be monitored by the WLC or WCS. They both provide several graphical reports. Active calls can also be monitored by CLI commands (see the examples below).

```
show client all
```

The optional *all* command shows all access points to which this client has associated. Information similar to the following is displayed:

```
AP Interface Mac: 00:0b:85:01:02:03
Client Interface Mac: 00:01:02:03:04:05
Measurement Duration: 90 seconds
Timestamp 21st Jan 2008, 06:35:80
UpLink Stats
=====
Average Delay (5sec intervals).....0
Delay less than 10 ms.....0
Delay bet 10 - 20 ms.....0
Delay bet 20 - 40 ms.....0
Delay greater than 40 ms.....0
Total packet Count.....0
Total packet lost count (5sec).....0
Maximum Lost Packet count(5sec).....0
Average Lost Packet count(5secs).....0
DownLink Stats
=====
Average Delay (5sec intervals).....15
Delay less than 10 ms.....1406
Delay bet 10 - 20 ms.....861
Delay bet 20 - 40 ms.....880
Delay greater than 40 ms.....363
Total packet Count.....4282
Total packet lost count (5sec).....772
Maximum Lost Packet count(5sec).....68
Average Lost Packet count(5secs).....2
Roam Count .....3
Roam Delay .....78
```

This information should be taken from an active call when at the desired data rate and transmit power when the client phone is at an anticipated cell edge. If the Average Lost Pocket count for five seconds is 0 or 1, then that link quality is ideal. This link should provide high-level MoS values if that area maintains low channel interference.

To see the CAC configuration for the 802.11a or 802.11b/g network, enter this command:

```
show {802.11a | show 802.11b}
```

To see the CAC statistics for a particular access point, enter this command:

```
show ap stats {802.11a | 802.11b} ap_name
```

```
Number Of Slots..... 2
AP Name..... AP.4e68
MAC Address..... 00:16:c7:d2:4e:68
Radio Type..... RADIO_TYPE_80211b/g
Stats Information
Number of Users..... 10
TxFragmentCount..... 1514948
MulticastTxFrameCnt..... 74
FailedCount..... 95714
RetryCount..... 647448
MultipleRetryCount..... 0
FrameDuplicateCount..... 92629
RtsSuccessCount..... 1889
RtsFailureCount..... 8874
AckFailureCount..... 638574
RxFragmentCount..... 0
MulticastRxFrameCnt..... 0
```

```

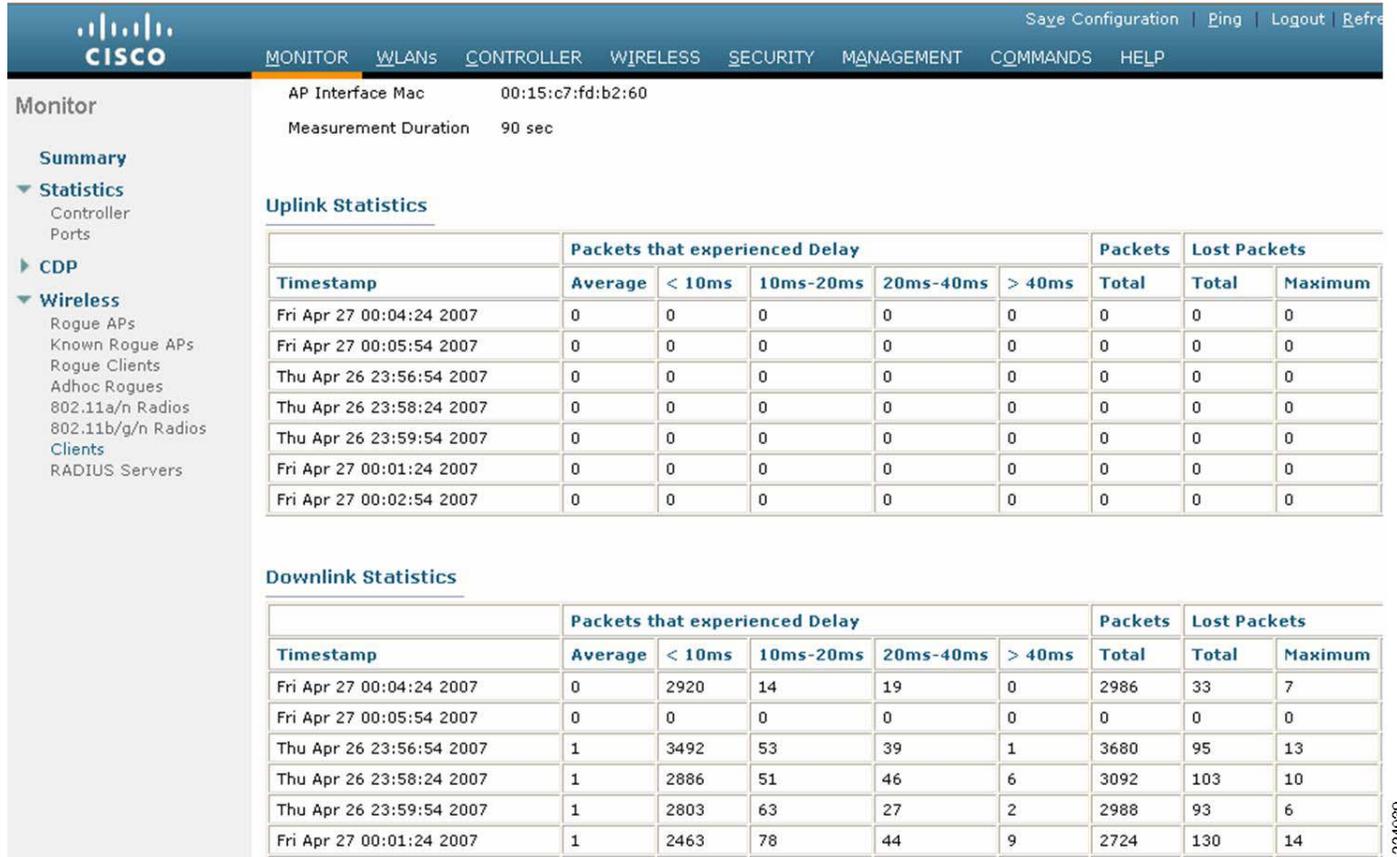
FcsErrorCount..... 927719
TxFrameCount..... 1774265
WepUndecryptableCount..... 0
TxFramesDropped..... 20943
Call Admission Control (CAC) Stats
Voice Bandwidth in use(% of config bw)..... 28
  Total channel MT free..... 0
  Total voice MT free..... 16656
  Na Direct..... 0
  Na Roam..... 0
Video Bandwidth in use(% of config bw)..... 0
Total num of voice calls in progress..... 3
Num of roaming voice calls in progress..... 0
Total Num of voice calls since AP joined..... 7
Total Num of roaming calls since AP joined..... 0
Total Num of exp bw requests received..... 0
Total Num of exp bw requests admitted..... 0
Num of voice calls rejected since AP joined.... 77
Num of roam calls rejected since AP joined.... 0
Num of calls rejected due to insufficient bw.... 77
Num of calls rejected due to invalid params.... 0
Num of calls rejected due to PHY rate..... 0
Num of calls rejected due to QoS policy..... 0

```

Knowing the counts of transmit and receive packets and retries is important to understanding the relative quality of the RF around this access point. Knowing the access point's historical performance may help determine if there is a new problem in its area. This can be useful in understanding the quantity of calls and the relative performance of calls on this access point.

The WLC provides GUI performance by VoWLAN client with download link and uplink packet delay statistics. As VoWLAN client drivers improve, there will be move and uplink statistics. [Figure 2-13](#) shows a CCX v4 client without uplink statistics.

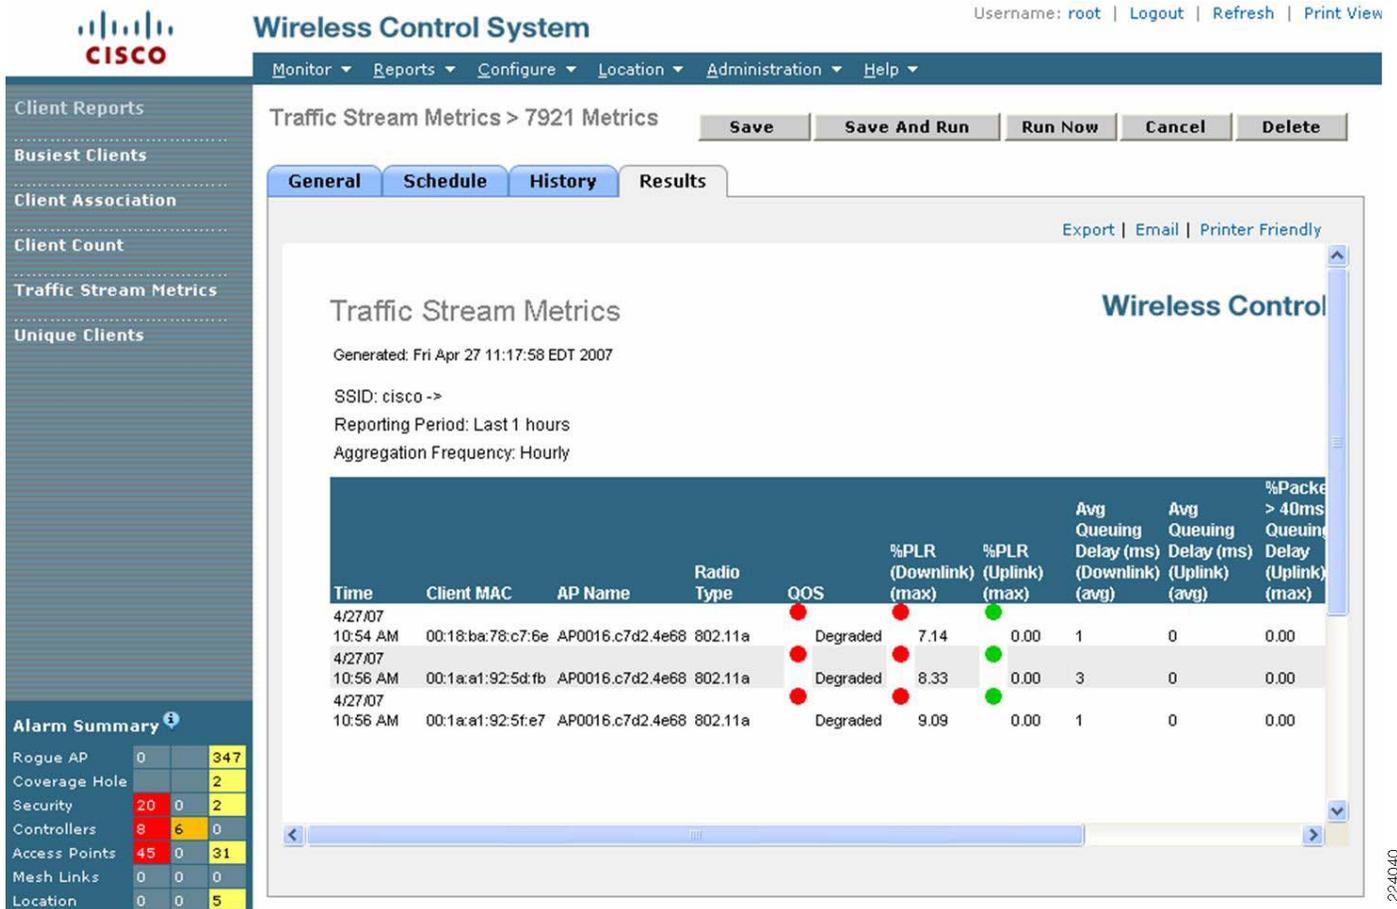
Figure 2-13 CCX v4 Client Without Uplink Statistics



224039

The WCS provides GUI by VoWLAN clients and access points. Figure 2-14 shows a report that was created by the WCS with a scheduled run time. The report can be exported and E-mailed.

Figure 2-14 Report Created with a Scheduled Run Time



These the statistics may vary for many reasons, such as amount of usage, where the VoWLAN client is used, and the utilization of channels during calls.

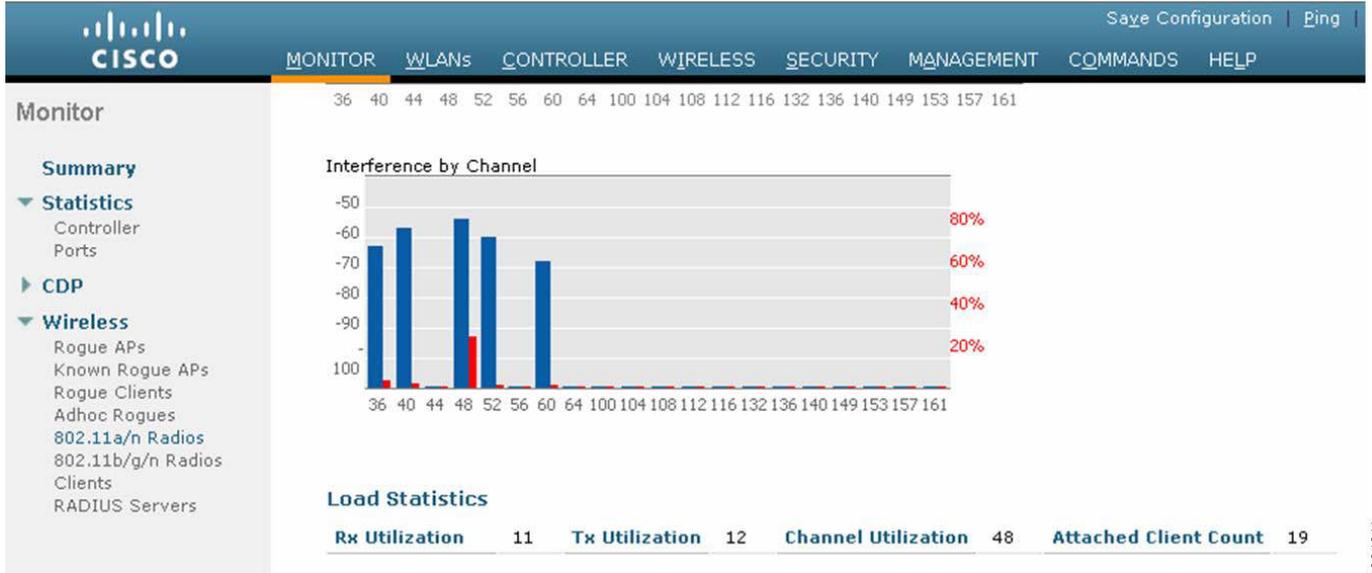
Establishing Baselines for Measuring and Managing WLAN Performance

It is important to establish a performance baseline early in the WLAN deployment. While it is never too late to establish a baseline, they should be established as soon as a new installation becomes stable. The baseline can be established from the GUI reporting of the WLC or the WCS. A VoWLAN baseline is just one of the performance baselines that should be established. There are many canned reports available on the WCS. The on-going management and continuing evaluation of WLAN performance must come from a well-defined baseline of network performance based on actual data collection, during real business hour testing, over an extended period of time. The baseline should exclude data that may have been gathered when there are interferers that will be removed from the site.

Channel interference is an important indicator of client performance. The more interference, the lower the signal to noise ratio (SNR). SNR is the measurement that is most often used to evaluate channel quality. Figure 2-15 shows the noise and interference on a 5GHz channel. Also shown at the bottom is the utilization of the access point's channel with the current client count. Note the noise levels reported

by the access points. There is a CCX option for clients to report noise to an access point and have that information forwarded in the background to the controller, but this is not common. For the most part the noise reported is data collected from the access points. This means there may be noise or interference that is near a client on the edge of the cell which would affect client performance, but that interferer may not be known to the access point.

Figure 2-15 Noise and Interference on 5GHz at a Site



224041

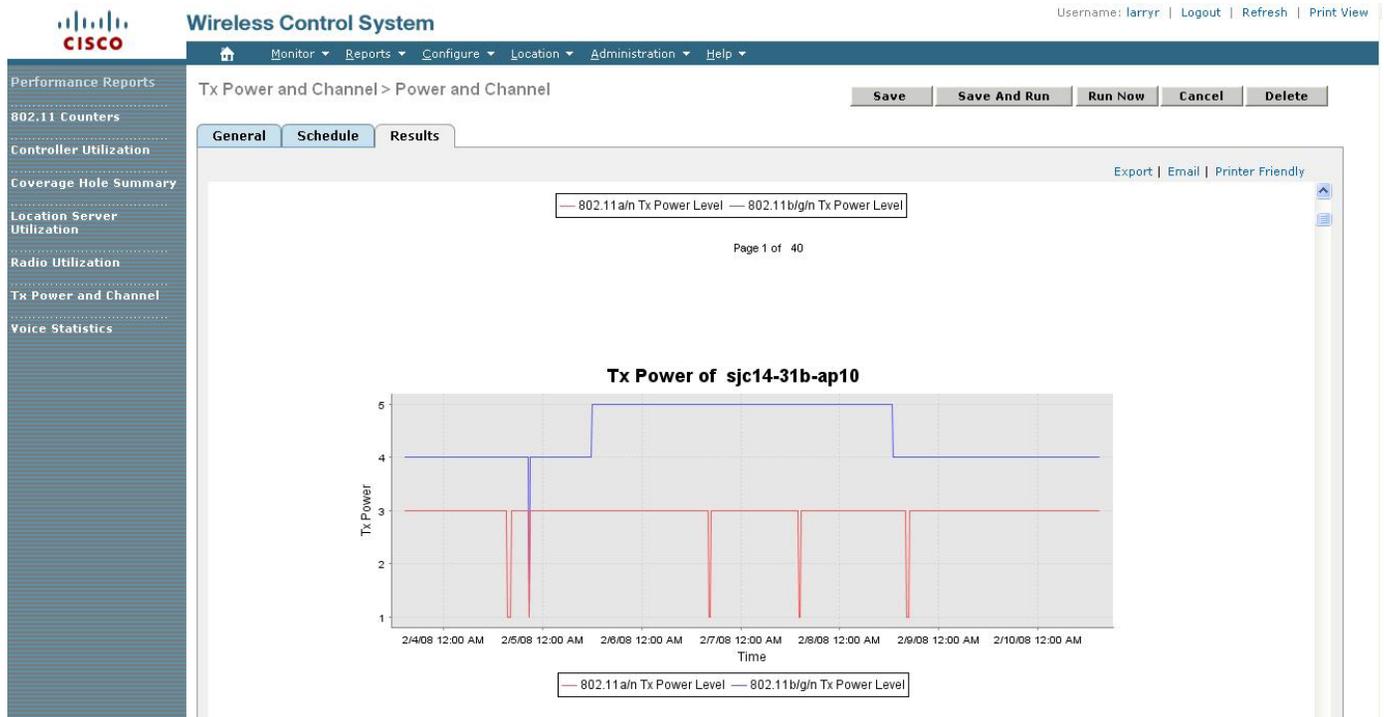
The noise on the channel is high with this high number of active clients. There are 19 clients on channel 48 in this test (see Figure 2-16).

Figure 2-16 Noise Profile



The WCS reports radio statistics by access point (see Figure 2-17).

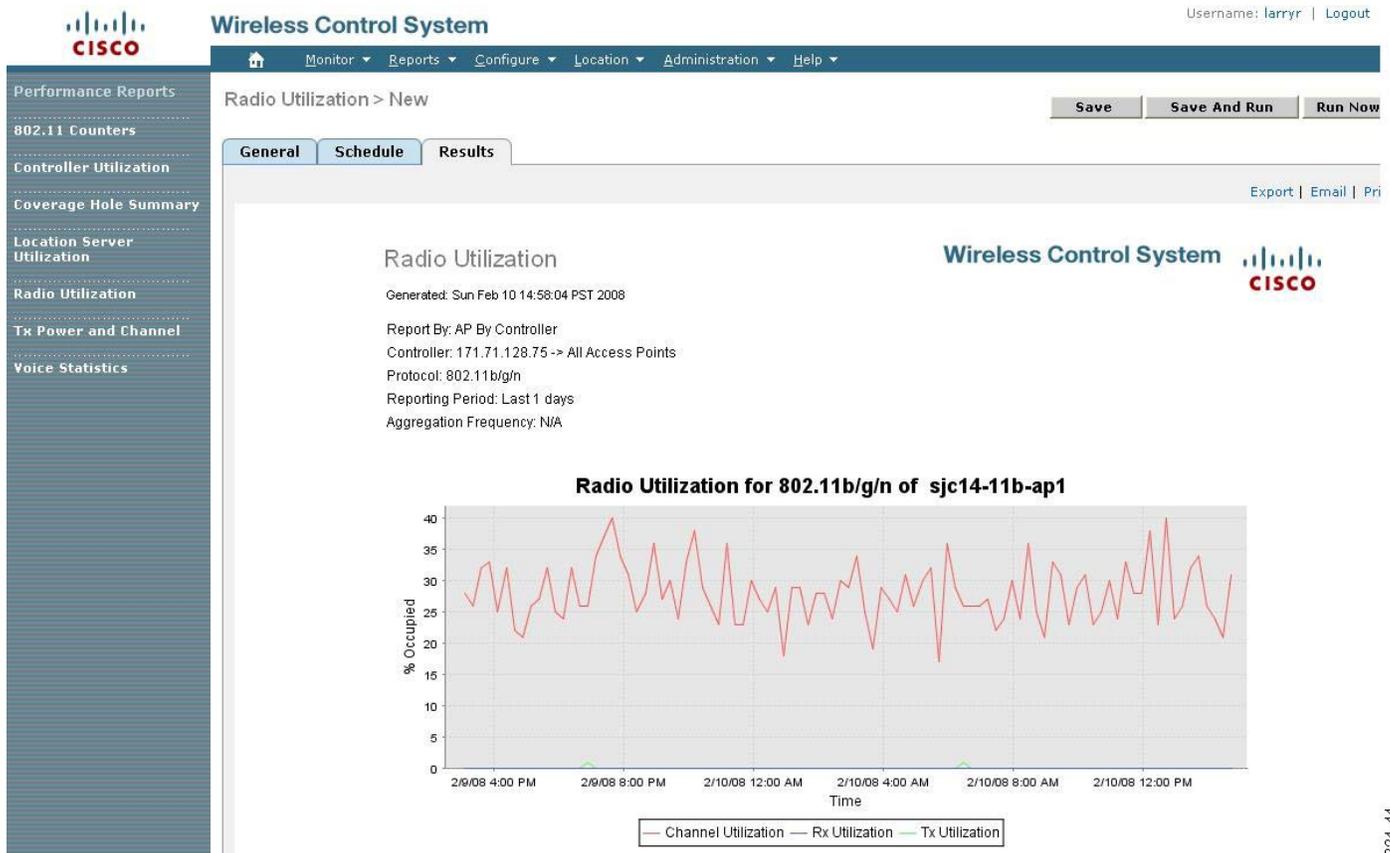
Figure 2-17 Radio Statistics by Access Point



224043

The WCS reports channels utilization statistics by access point (see Figure 2-18).

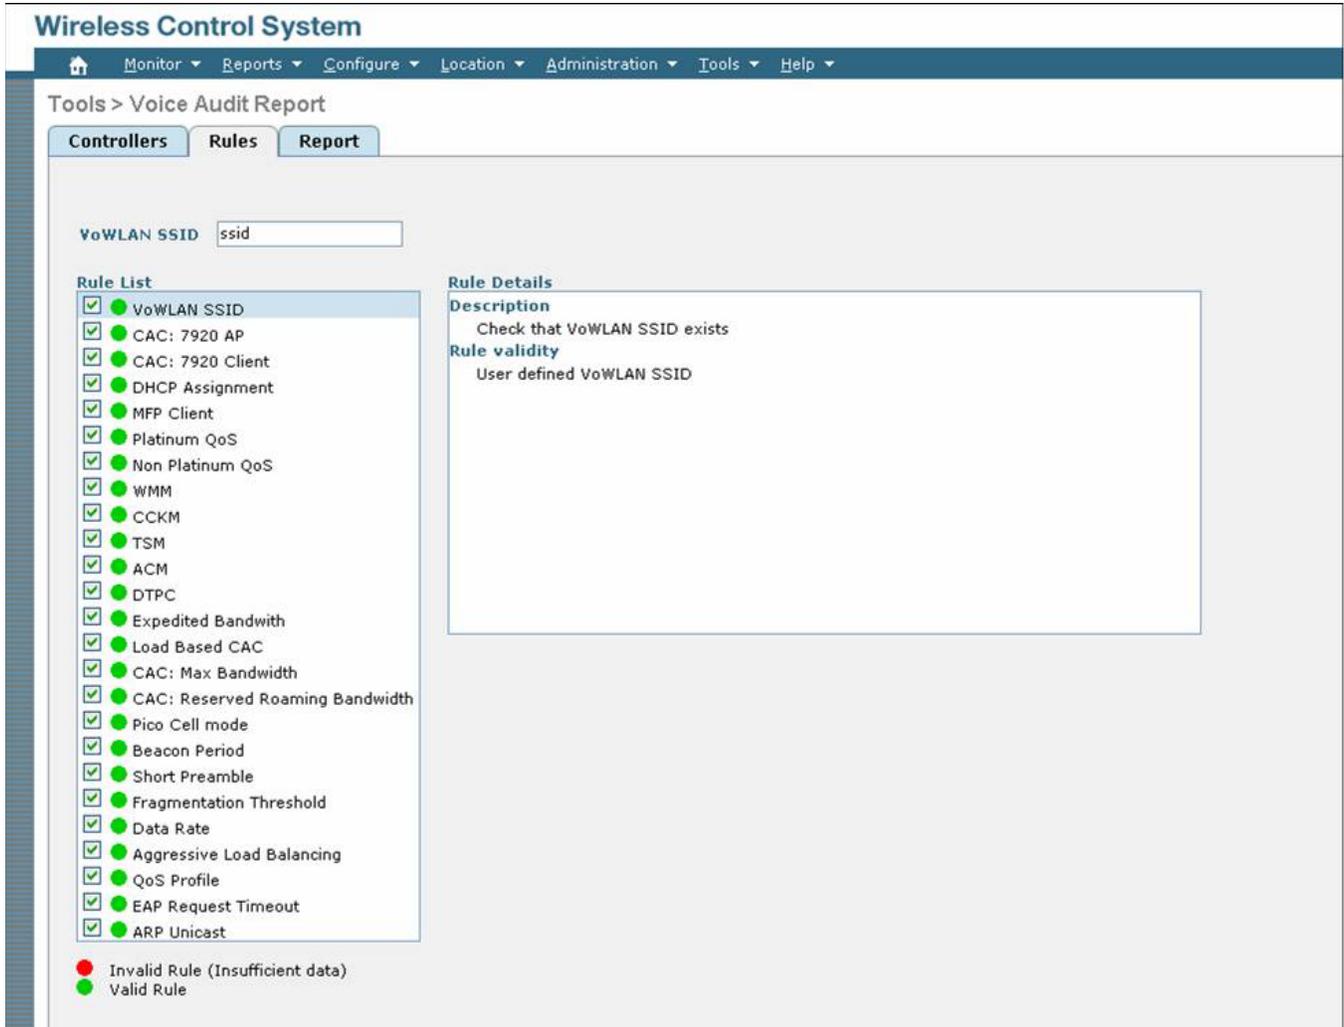
Figure 2-18 Channels Utilization Statistics by Access Point



The configuration on controllers changes over time. Those configurations should also undergo periodic audits. The WCS audits the configurations of selected controllers or all the controllers reachable by the WCS.

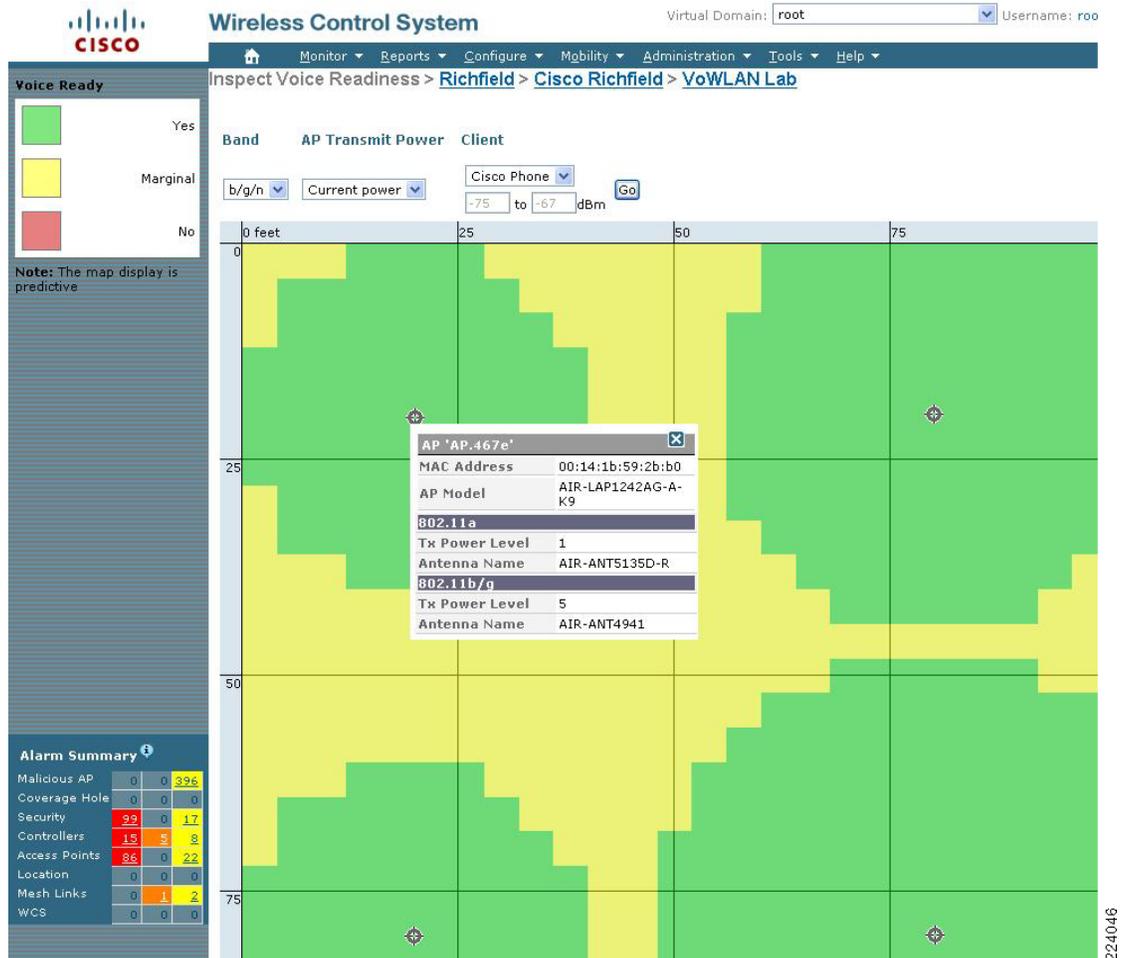
As the VoWLAN configuration was initially audited during the planning stage, it should be periodically audited (see Figure 2-19).

Figure 2-19 Voice Audit Report



The floor coverage should also be audited to see if all cells still provide the necessary coverage levels (see Figure 2-20).

Figure 2-20 Floor Coverage Audit



The WCS also provides numerous client statistics and access point statistics that provide insight into the current behavior of individual clients and their interactions with the access points. The reporting can be used to see how often a client roamed, identify the access point to which it roamed, and the times involved.

show client tsm

The optional *all* command shows all access points to which this client has associated. Information similar to the following is displayed:

```
AP Interface Mac: 00:0b:85:01:02:03
Client Interface Mac: 00:01:02:03:04:05
Measurement Duration: 90 seconds
Timestamp 21st Jan 2008, 06:35:80
UpLink Stats
=====
Average Delay (5sec intervals).....0
Delay less than 10 ms.....0
Delay bet 10 - 20 ms.....0
Delay bet 20 - 40 ms.....0
```

```

Delay greater than 40 ms.....0
Total packet Count.....0
Total packet lost count (5sec).....0
Maximum Lost Packet count(5sec).....0
Average Lost Packet count(5secs).....0
DownLink Stats
=====
Average Delay (5sec intervals).....15
Delay less than 10 ms.....1406
Delay bet 10 - 20 ms.....561
Delay bet 20 - 40 ms.....880
Delay greater than 40 ms.....363
Total packet Count.....4282
Total packet lost count (5sec).....772
Maximum Lost Packet count(5sec).....68
Average Lost Packet count(5secs).....0
Average Lost Packet count(5secs).....0
Roam Count.....3
Roam Delay.....78

```

This information should be taken from an active call when at the desired data rate and transmit power when the client phone is at anticipated cell edge. If the Average Lost Packet count for five seconds is 0 or 1, then that link quality is ideal. This link should provide high level MoS values if that area maintains low channel interference.

To see the CAC configuration for the 802.11a or 802.11b/g network, enter this command:

```
show {802.11a | show 802.11b}
```

To see the CAC statistics for a particular access point, enter this command:

```
show ap stats {802.11a | 802.11b} ap_name
```

```

Number Of Slots..... 2
AP Name..... AP.4e68
MAC Address..... 00:16:c7:d2:4e:68
Radio Type..... RADIO_TYPE_80211b/g
Stats Information
  Number of Users..... 10
  TxFragmentCount..... 1514948
  MulticastTxFrameCnt..... 74
  FailedCount..... 95714
  RetryCount..... 647448
  MultipleRetryCount..... 0
  FrameDuplicateCount..... 92629
  RtsSuccessCount..... 1889
  RtsFailureCount..... 8874
  AckFailureCount..... 638574
  RxFragmentCount..... 0
  MulticastRxFrameCnt..... 0
  FcsErrorCount..... 927719
  TxFrameCount..... 1774265
  WepUndecryptableCount..... 0
  TxFramesDropped..... 20943
Call Admission Control (CAC) Stats
  Voice Bandwidth in use(% of config bw)..... 28
    Total channel MT free..... 0
    Total voice MT free..... 16656
    Na Direct..... 0
    Na Roam..... 0
  Video Bandwidth in use(% of config bw)..... 0
  Total num of voice calls in progress..... 3
  Num of roaming voice calls in progress..... 0
  Total Num of voice calls since AP joined..... 7

```

```

Total Num of roaming calls since AP joined..... 0
Total Num of exp bw requests received..... 0
Total Num of exp bw requests admitted..... 0
Num of voice calls rejected since AP joined.... 77
Num of roam calls rejected since AP joined..... 0
Num of calls rejected due to insufficient bw.... 77
Num of calls rejected due to invalid params.... 0
Num of calls rejected due to PHY rate..... 0
Num of calls rejected due to QoS policy..... 0

```

Knowing the counts transmit and receive packets and retries is important to understanding the relative quality of the RF around this access point. Knowing the access point's historical performance may help determine if there is a new problem in its area. This can be useful in understanding the quantity of calls and the relative performance of calls on this access point.

Location Based Services (LBS) should be baselined by putting reference tags in specific locations and documenting their locations after doing Point Mode and Linear Mode Calibrations. Figure 2-21 shows the Point Mode and Figure 2-22 shows the Linear Mode Calibrations.

Figure 2-21 Point Mode Calibration

The screenshot shows the Cisco Wireless Control System interface. The main area displays the calibration process for model 'EW-10/08/07' on Floor Nortech > Amber. The collection method is set to 'Point' with a zoom of 100%. A floor plan is shown with a yellow highlighted area for calibration. A legend on the left indicates coverage for 802.11a/n, 802.11b/g/n, and 802.11a/b/g/n. An alarm summary table is also visible.

Alarm Summary			
Rogue AP	0	0	307
Coverage Hole	0	0	12217
Security	8	0	3
Controllers	0	0	0
Access Points	0	0	15
Location	0	0	0
Mesh Links	0	0	0

224047

Figure 2-22 Linear Mode Calibration



Either of these two methods can be used to establish the baseline for an area of the hospital. More details on the calibration process can be found in the WCS design guide. The location accuracy tool can then be scheduled to run at regular intervals or on demand. The reference tags are used to generate detailed reports on accuracy over a period of time.

There are many reporting options with the WLC or WCS. that can be used for maintenance and trouble shooting.

Considerations for Alternate or Third-Party Antenna Systems

Prior to deploying Distributed Antenna Systems (DAS) or leaky coax systems, it is necessary to determine how such alternate antennas systems affect the collision domain, cell size, cell throughput, and the performance of your applications. It is necessary to know what frequencies, antenna types, and 802.11 radios they support.

When deploying DAS, the access points are centralized in a single location or co-located in a single enclosure or equipment closet. Unfortunately, aggregating access points in one location and using distributed antennas violates several of the technical requirements for optimal WLAN functionality and coverage. When multiple antennas are connected to a single access point, degradation or impairment of most of Cisco's advanced features in the Cisco Unified Wireless Network should be expected because it is impossible to reuse channels to maximize throughput when the access points are aggregated and the channels are shared across the entire antenna system.

In active DAS systems, passive antennas are at the end of regular coax. In passive DAS systems, leaky cable is distributed around the floor. When deploying DAS with the Cisco Unified Wireless Network, the signal strength also decreases and creates poor application performance. The use of a radiating cable or a leaky coaxial type of deployment also leads to performance degradation because this introduces loss

in the antenna system on both the transmitter and the receiver. To avoid this situation, Cisco recommends the use of one diversity antenna to a single access point radio. And in the case of an 802.11n radio, the antenna needs to be a MIMO antenna solution. This is necessary to provide better wireless coverage and application performance. Outside of this recommended design, only best effort Cisco TAC support can be provided to customers.

The Cisco WLAN controllers enable Radio Resource Management (RRM), which provides advanced and automated management capabilities and enhanced performance. RRM allows Cisco's Unified WLAN Architecture to continuously analyze the existing RF environment, automatically adjusting the power levels of access points and channel configurations to help mitigate channel interference and signal coverage problems. RRM also helps to increase system capacity and provide automated self-healing functionality to compensate for RF dead zones and AP failures.

RRM and the Cisco Location solution are designed around antennas with well-defined RF characteristics. The addition of any third-party antenna results in non-standard (i.e., not tested or validated) results. Some RRM adjustments may be required to achieve optimal performance, especially for customers who have deployed the Cisco Location Solution or VoWLAN services. For more technical information, refer to the Cisco Application Bulletin that highlights the limitations of deploying DAS with the Cisco Unified Wireless Network.

The motivations for a DAS architecture and the main benefit is the ability to offer services via multiple, discrete wireless systems for more pervasive coverage. Healthcare providers deploy DAS solutions to manage several wireless services, including clinical data, computerized physician order entry (CPOE), paging, public cellular services, nurse-call systems, patient telemetry, and future monitoring applications. The system can also distribute two-way radio for security and maintenance, as well as Wi-Fi and cellular phone access for patients, visitors, and staff. Hospitals are also looking to centralize management of these wireless systems. Hospitals are also delivering Wi-Fi from the DAS architecture to limit ceiling penetration as they may have infection control mechanisms in place, making access above the ceiling tiles cost prohibitive.

Benefits

- **Convergence**—Ability to carry multiple wireless access technologies, such as cellular and Wi-Fi, simultaneously and allow the building planner to provide coverage for many diverse services, such as cellular phone, two-way radios, paging, Wi-Fi, Medical Telemetry systems, etc.
- **Time savings**—IT managers can deploy both wireless technologies with a single cabling at the same time.
- **Physical security for network equipment**—The access points are secured in a locked wiring closet rather than in ceilings or walls near users' desks, reducing the likelihood of damage or theft.
- **Network maintenance**—Easier configuration, maintenance, and replacement of access points since they are located in a single, easily accessible location.

Limitations

- **Unpredictable performance**—Performance, such as capacity and throughput, are highly dependent on site design and installation and output power. Connecting antennas to aggregated access points with coax cabling significantly reduces the output power at each antenna. For example, a long coax cable that runs between access points and antennas introduces loss in the antenna system on both the transmitter and the receiver.

- Performance degradation for 802.11a and 802.11n—Performance is degraded network wide, however the edge of the network may provide acceptable performance because it is impossible to reuse channels to maximize throughput when the access points are aggregated and the channels are shared across the entire antenna system. The reduced output problem is worse with 802.11a or 802.11n installations because the 5GHz frequency attenuates much more quickly. The signal strength to transmit large files or support voice quality of service may also be lower.
- Performance degradation for MIMO 802.11n—The 802.11n standard uses multiple antennas with MIMO technology. It is unclear how newer technologies that rely on multiple antennas can be leveraged with a DAS system since this limits user capacity. There is currently no upgrade path to support 802.11n because most DAS solutions are based on one antenna. Also, the high attenuation caused by the coax DAS system and distortion issues created by aggregating access points makes it difficult to provide even acceptable capacity.
- Radio Resource Management—RRM may not perform as expected due to the change in access point site design and increased network complexity. Some RRM adjustments need to be made to achieve optimal results.
- Location—A thorough network design and site survey must be completed since aggregating access points in one location decreases the location accuracy. Some RRM adjustments need to be made to achieve optimal results.
- Voice—Voice QoS may be lower due to lower signal strength and lack of diversity antenna support. VoWLAN deployments also require a denser population of access points than most data-oriented application deployments and therefore must be designed properly to maintain QoS and minimize delay and jitter.

Cisco Technical Assistance Center (TAC) will make a best effort to support Cisco Unified Wireless Network deployed along with a third-party DAS solution. Indeed, Cisco TAC does not have access to data on these third parties' products and their quality issues, manufacturing defects, support history, etc. and cannot easily replicate such a system. In the case of RF coverage related problems, Cisco particularly encourages customers to contact the system integrator who installed the DAS as design and installation are crucial in defining the performance of the network. Customers could experience a longer problem identification and isolation, which could lead to longer resolution times for RF and coverage problems.

VoWLAN and Location services are applications that perform best when cell sizes are small and multiple antennas provide signal reception to the access point radios.



CHAPTER 3

Common Wireless Medical Devices

Overview

This chapter provides an overview of the wireless 802.11-based devices used in medical facilities. Both wired and wireless devices are used throughout the healthcare enterprise, therefore aggregate traffic load and network metrics must be considered in the overall wireless network design. The trend within healthcare continues to be the rapid growth of wireless-enabled medical devices with various traffic requirements. This document focuses on 802.11-enabled devices, their network characteristics, wireless considerations, and several other factors are described for each device type. In order to discuss medical device types, a baseline categorization of the various devices is important as it provides a framework for their use in a hospital environment.

While there are many ways to classify the use of wireless devices, including classification models that customers may develop, the two sample classification models described in this section are intended to stimulate discussion about how device classification provides a model for monitoring and designing a wireless network for specific device types.

- **Medical devices**—These devices are directly associated with patient care. Applications may be included in this category if the application information is used directly for patient care. Typically, examples of these devices include infusion pumps, patient monitors, ventilators, and other similar devices. Clinical applications, with the exception to EMR systems which are considered collectors of clinical information, also fall into this category.
- **Components**—Integral parts of devices which in some instances may be viewed as devices. Special monitors, probes, or similar may fall into this category.
- **Accessories**—Additional devices and applications which assist in patient care and are linked to a device.

Another example of medical device classification would be through their use model or role, as shown in [Table 3-1](#).

Table 3-1 **Role-Based Medical Device Classification**

Priority	Description
Highest	Examples include devices that are used by hospital staff that are life critical to the care of a patient.
High	These could be devices used by hospital staff that aide in their day-to-day workflow, but are not life critical.
Medium	Devices used by hospital staff for business administration is one example.
Low	Devices used by guest or patients.

These two sample classification models are intended to illustrate ways of grouping devices that are filtered by answers to a few qualification questions. While this list of questions is not inclusive of all aspects of device classification that a healthcare provider must consider, it is presented here as a baseline for consideration:

- What is the importance of a particular device as it matches against the business metrics and patient care for that hospital facility?
- What service level is required by the network to deliver the services for that device (such as QoS, uptime and availability, regulations, etc.)?
- What type of wireless coverage is required and has the proper RF site survey been performed to ensure adherence to the coverage requirements for the metrics defined?
- What applications are used on the medical devices that are wireless enabled?
- Do any of the applications require unique network considerations that may conflict with typical network design best practices?
- What features do the clients support to ensure compatibility with the wireless infrastructure (examples include the security features and QoS capabilities that the endpoints support to inter-work with the wireless system)?

To plan an initial build or expansion of a healthcare environment's wireless network, the medical facility should, in addition to performing a complete RF site survey, inventory the wireless-enabled devices that are currently used or planned to be used. These are then used as qualification questions (adding others as appropriate) to create a categorization model that best meets the unique requirements of the healthcare organization.

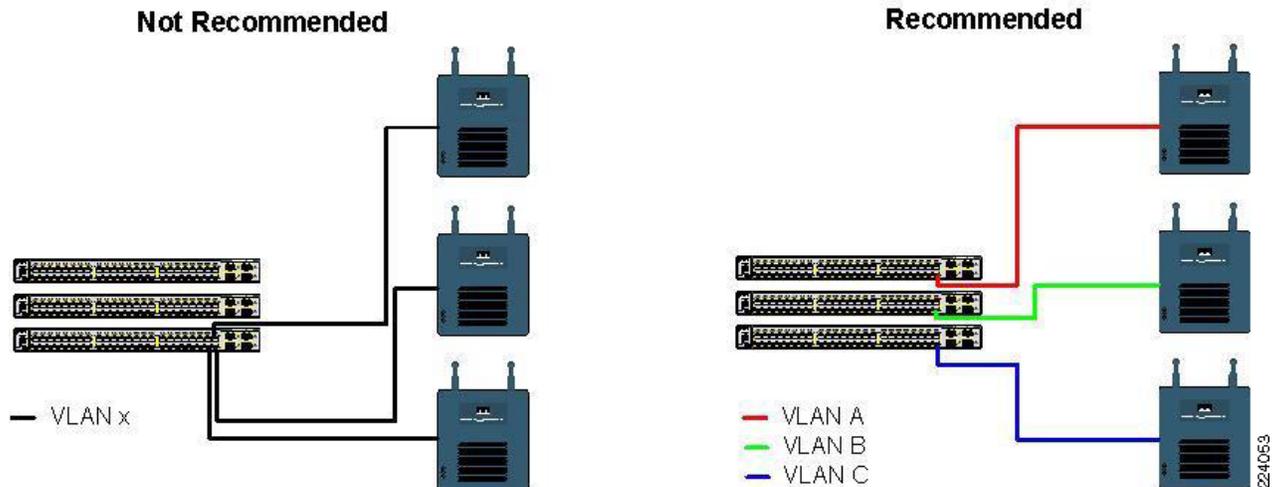
In most wireless installations, the wireless network is built on top of the existing wired Ethernet network. As such, the reliability of the wireless infrastructure can only at best mirror that of the reliability of the wired network. It is therefore recommended to perform a gap analysis of the overall wired network infrastructure as it relates to high availability (HA). Some common items to include in the gap analysis include:

- Redundant switches or switch fabric at the access layer
- Redundant uplinks from the access layer to the distribution or core
- Ensure that two or more diverse and redundant power feeds to the access layer and distribution layers
- Current and available inventory of cold spare network hardware for access, distribution, and core infrastructure
- A well-planned-for and implemented end-to-end QoS policy
- A tested and functional Layer 3 routing protocol which provides proven rapid convergence
- Elimination of outdated Layer 2 loop control and the adoption of optimized Layer 2 loop control spanning tree protocols such as Rapid per VLAN Spanning Tree (RPVST)
- 802.3af-capable switches supporting Class III power (15.4 Watts) for non-802.11n deployments and PoE+ for those environments where 802.11n is required
- Availability of ports, and potentially line card slots, for the redundant Wireless LAN Controllers (WLC) within the distribution layer

Once corrective measures have been applied to the existing wired infrastructure, the overall network should be validated for rapid convergence in the event of component or uplink failure. Once completed, the installation of the wireless infrastructure can begin. An often overlooked design aspect during the deployment of a wireless network is the manner in which the access points are connected to the access layer. Frequently ports are reserved at the end of the switch stack or chassis for access points. When the

installation is complete, all of the access points serving a floor are connected to a single switch promoting a single point of failure, therefore nullifying all of the considerations and measures that were implemented to provide the upstream redundancy. It is therefore highly recommended to spread the access points across multiple switches or line cards to eliminate all single points of failure.

Figure 3-1 Access Point Connections to Access Layer



As shown in [Figure 3-1](#), separating the access-points across a number of different access switches, or in some cases line cards within a chassis, helps eliminate a single point of failure. In addition, and if possible, the APs can also be spread across separate VLANs if present within the access layer. This further enhances the redundancy design by reducing the likelihood of an issue specific to a single VLAN impacting all access points on a given floor. Although not shown in [Figure 3-1](#), each AP on a patient floor should be serviced by the same Wireless LAN Controller (WLC) to provide optimum roaming between access points and thus prevent unnecessary roaming between controllers.

In summary, to ensure that the wireless network can provide a high level of availability, the underlying wired infrastructure on which the wireless network is built must be evaluated. Cisco recommends that the overall characteristics of an HA design be implemented end-to-end to provide a highly-available wireless design. Best practices for Wireless and Campus designs can be found at www.cisco.com/go/cvd and www.cisco.com/go/srnd.

Wireless QoS Considerations

Wi-Fi Multimedia (WMM) is a subset of the 802.11e specification which has been certified by the Wi-Fi Alliance for interoperability. The Wi-Fi Alliance issues WMM certifications for devices that are submitted and pass the Wi-Fi WMM tests. The Cisco Unified Wireless Controllers and Cisco Aironet access points are part of the Wi-Fi Alliance test bed, which means all client devices that have the WMM certifications were directly tested against WMM-configured Cisco Wireless LAN Controllers.

The original 802.11 specifications did not include QoS or call management. If wireless clients in the same cell transmitted simultaneously, a packet collision occurs. To avoid collisions and therefore communicate effectively, wireless clients must take turns transmitting. The original 802.11 standard defined two methods for sharing access to the RF channel:

- Distributed Coordination Function (DCF)—Stations using DCF monitor the channel to determine if another device was transmitting.

- Point Coordination Function (PCF)—A polling technique in which the access point assigns the clients periods of access time slots.

DCF did a good job of letting stations send data while avoiding collisions, but it does not provide a mechanism for assigning priority in the cell or on the RF channels and is therefore not deterministic in nature. Before WMM, there was no way to prioritize access to the RF channel or to maintain channel utilization. Today's wireless networks still require a WMM-aware client that adheres to the WMM specification to prevent the channel from being over-utilized and prevent VoWLAN calls from saturating a wireless cell or service area.

VoWLAN clients that are WMM certificated and are associated to an SSID configured on the controller to support the WMM specification have QoS priority over the air. The WMM specification defines eight user priorities. There are four access categories defined by 802.11e.

The 802.11e Enhanced Distributed Channel Access (EDCA) mechanism uses 802.1d user priority (DiffServ tags) to classify the traffic categories:

- Voice—Priority 7 or 6 for toll-quality VoWLAN calls requiring low latency.
- Video—Priority 5 or 4 for SDTV or HDTV video streams.
- Best Effort—Priority 3 or 0 for latency-insensitive, interactive applications.
- Background—Priority 2 or 1 for batch data transfer applications.

Clients that do not support the WMM standard are referred to in the 802.11e specification as non-QoS stations. **Data transmitted by non-QoS stations is classified as best effort, and therefore do not receive any over-the-air prioritization regardless of the configuration of the SSID to which they are associated.**

Clients not able to support WMM QoS, because of firmware or driver limitations, can be assigned to SSIDs that service wireless clients capable of supporting WMM. The performance of a cell supporting both WMM capable and non-capable wireless clients is still enhanced to a degree. This is because the traffic sent in the direction of the access point to the wireless client are in one of the four access categories and marked with one of the eight user priorities. See [Figure A-5 in Appendix A, “Antenna Recommendations”](#) to see the marking of a VoWLAN transiting from the VoWLAN client to the WLAPP controller and back to the phone.

Figure 3-2 WMM Policy Setting for the 802.11e Access Category for Voice



In [Figure 3-2](#), the WMM policy setting for the 802.11e access category for voice would mark the voice packets going to the VoWLAN clients associated to SSID for voice WLAN with a user priority of six. The three other categories are selectable from the drop-down QoS field shown above.

The Use of Unlicensed 802.11 Frequencies in Healthcare

There are three unlicensed bands at 900 MHz, 2.4 GHz, and 5.7 GHz. These bands are referred to as the Industrial, Scientific, and Medical (ISM) frequencies. The Federal Communications Commission (FCC) regulates the use of these bands in the US. The FCC regulates the transmit power in these bands and how that power is modulated. Before a radio can be used in the United States, it must pass a set of rigid FCC tests. The FCC also rules on antennas and radio/antenna combinations as it relates to effective radiated power.

The Institute of Electrical and Electronics Engineers (IEEE) maintains the 802.11 specifications. The 802.11 radio-related specifications follow the FCC regulations. Current 802.11 specifications utilize the 2.4GHz and 5GHz ISM frequencies. The ISM frequencies are unlicensed, which means that the user or organization using the spectrum does not have to register with the FCC as to the frequencies being used and the physical location(s) they are being deployed. The use of devices in the ISM frequencies is defined by Part 15 of the FCC regulation. Unlike licensed frequencies, for example local radio or television stations, there are no distance separation requirements in the ISM frequency bands.

The excerpt below is from a paper written by N. Golmie of the National Institute of Standards and Technology and is an excellent summary of the interference issues in the 2.4 GHz Band. The 2.4 GHz ISM band allows for primary and secondary uses. Secondary uses are unlicensed, but must follow rules defined in the Federal Communications Commission Title 47 of the Code of Federal Regulations Part 15 [COM] relating to total radiated power and the use of the spread spectrum modulation schemes. Interference among the various uses is not addressed as long as the rules are followed. Thus, the major downside of the unlicensed ISM band is that frequencies must be shared and potential interference tolerated by all wireless clients using the band.

The use of spread spectrum and its associated power rules are fairly effective in dealing with multiple users in the band, provided the radios are physically separated. The same however cannot be said for close proximity radios. Multiple users, including self-interference of multiple users of the same application, have the effect of raising the noise floor in the band, resulting in a degradation of performance. The impact of interference may be even more severe when radios of different applications use the same band while located in close proximity.

Interference in the 2.4 GHz ISM Band:

Challenges and Solutions

N. Golmie

National Institute of Standards and Technology

Gaithersburg, Maryland 20899

Part 15 of the FCC code also includes the regulations regarding antennas; transmit powers of radios and how the antennas and radios are used together in the ISM bands.

The Effective Isotropic Radiated Power (EIRP) of a transmitter is the power that the transmitter appears to have if the transmitter were an isotropic radiator (if the antenna radiated equally in all directions).

By virtue of the gain of a radio antenna (or dish), a beam is formed that preferentially transmits the energy in one direction. The EIRP is estimated by adding the gain (of the antenna) and the transmitter power (of the radio).

$EIRP = \text{transmitter power} + \text{antenna gain} - \text{cable loss}$

When using radio equipment, there are limits on the output of the system. These limits are given as EIRP, and must not be exceeded. Different countries will have different standards.

Check with authorities in the country of installation to determine maximum EIRP. The maximum EIRP allowed by the FCC for a Part 15 802.11b device in the United States is 36 dBm.

The standards are different for specific point-to-point systems. Indoor WLANs that would be considered point-to-multipoint solutions, so the maximum EIRP allowed must not exceed 36 dBm and the maximum gain on an antenna must not exceed 16 dBi (for the United States) unless installed by a professional installer. A 30dBm maximum transmitter power with 6dBi maximum antenna gain for an EIRP value of 36 dBm is a FCC rule. The Regulatory Domains of the other areas of the world like ETSI in Europe have similar rules.

Other government bodies may have additional rules. The current ETSI standards are EN 300 328-1 V1.2.2. The documentation is available from American National Standards Institute (ANSI).

The FCC rules for 5GHz are different than their rules for 2.4GHz. The EIRP rules actually change depending on the frequency band. The UNII-1 channel 36 has a maximum FCC transmit power of 17 dBm while it is 30 dBm for channel 52. In Singapore the maximum transmit power of channel 36 is 13dBm. To stay compliant order the product part number that matches your regulatory domain. And have the controller configured to the correct domain. The guide to correct regulatory domains is listed in [Appendix B, "References"](#) with a link to download it.

RF Interference Common to Healthcare environments

Because the 802.11 RF spectrums (2.4 GHz and 5 GHz) are both unlicensed, any device operating within these bands must tolerate interference from other non-licensed devices. **Such interference is highly likely and should therefore be planned for and expected in all installations.** In the 2.4GHz ISM band which is used for 802.11b/g, the likelihood of interference is currently higher than that of the 5 GHz 802.11a band due to the large number of consumer devices sharing the band. Some sources of interference are transient in nature, making their detection and isolation difficult.

One of the most common devices found in a healthcare facility that uses the 2.4GHz band is the microwave oven. The maximum output power of an 802.11b/g access point is typically 1/10 of a watt or 100 milliwatts. Microwave ovens, by comparison, typically operate between 700 and 1400 watts and are centered at 2.45GHz, centered at about two thirds of the way in the ISM 2.4GHz band. This higher power equates to 14,000 times the maximum power allowable for 802.11b/g devices operating in the 2.4 GHz band. The good news is that microwave ovens do have RF shielding to prevent the radiation of RF energy beyond the inside of the oven. The bad news is that some poorly-designed or malfunctioning ovens may in fact radiate some of this unwanted RF energy into the surrounding area.

If such an oven is found, it should be removed from service and replaced with a properly operating oven. In the event that a microwave oven is causing interference in the 2.4GHz band, most of the interference will be experienced on 802.11b/g channels 6 and 11. This is because a microwave oven operates at 2450 MHz and shifts frequency up and down slightly during operation. By using channel 1, which resides between 2401MHz and 2423 MHz, the interference effects of a microwave oven can be avoided. Another approach would be to avoid the 802.11b/g 2.4GHz band entirely and use the 802.11a 5GHz band instead. More specific information about RF radiation standards as they relate to microwave appliances can be found on the FDA Web site at: <http://www.cfsan.fda.gov/~lrd/FCF1030.html>.

Examples of other devices that also operate in the 2.4 and 5 GHz ISM band(s) are:

- Bluetooth devices (Frequency Hopping, FH)
- 2.4 GHz and 5.8 GHz cordless phones (Frequency Hopping, FH)
- 2.4 GHz wireless video systems and video senders (non-802.11 based)
- Older 2.4 GHz FH bar code scanners
- Remote-controlled toys (cars, planes, robots etc.)
- 5.8 GHz door motion detectors

- Wireless USB devices
- Some forms of aviation-based radar systems
- Some wireless headphones
- Dual-mode or Wi-Fi equipped cell phones

A wireless network that is implemented and operating perfectly today could tomorrow be unusable due to new devices causing interference. Because a hospital can be generally thought of as a public space, the introduction of other consumer-class devices generating interference is also quite common. It is therefore recommended that backup mechanisms be implemented to provide like service over a wired infrastructure in the event of a disruption to the wireless network.

Since it is impossible to predict and prevent these unwanted interference sources, an RF baseline is one technique that can be used to determine when the wireless network is encountering interference. Baselineing the wireless RF spectrum should be performed during the implementation and turn-up phase of all wireless networks. The baseline document notes the noise floor and signal quality for multiple locations on each floor and includes patient rooms on both sides of a corridor, not just on one side. The information can then be used when a problem is encountered and interference is suspected. Without the baseline, it's not possible to determine what is normal or abnormal for any given RF environment.

Quick detection and isolation of interference sources within a healthcare environment is critical to preventing workflow disruptions. A spectrum analyzer is a tool that can graphically show you the RF fingerprint of a range of frequencies or band. Through the use of Cisco's Spectrum Expert tool (formerly Cognio), the process of determining the type of device causing the interference is automated. The Spectrum Expert tool has the ability to classify sources of RF interference by type and manufacturer. The tool can quickly locate and isolate many types of RF interference in both the 2.4 GHz and 5 GHz bands.

As shown in [Figure 3-3](#), each RF device has a particular fingerprint. For non-802.11-based sources of interference, the fingerprint includes the frequency or frequencies used, duty cycle, frequency shifts, and duration. In all cases, this fingerprint is used to determine the type of device (DECT phone, Bluetooth, microwave, video sender, etc.) and possibly even the manufacturer. By using this information and directional antennas, the offending device and its physical location can be quickly determined, reducing the disruption to the clinical workflow. [Figure 3-3](#) and [Figure 3-4](#) provide some examples of automatically classified sources of interference, specifically a bluetooth device and a microwave oven.

Figure 3-3 Automatically Classified Source of Interference—Bluetooth Device

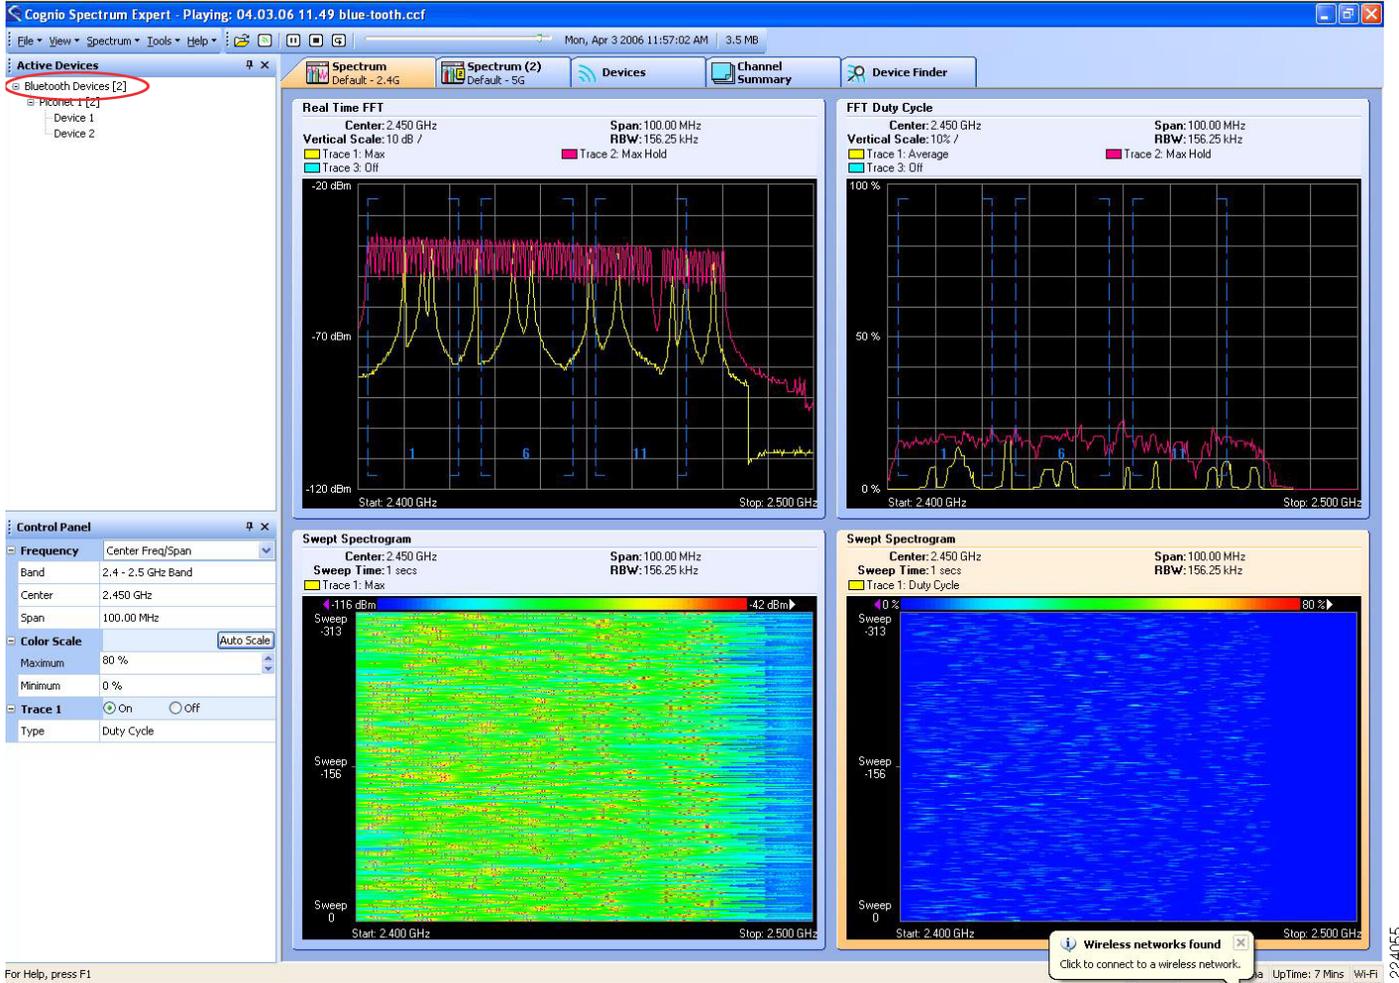
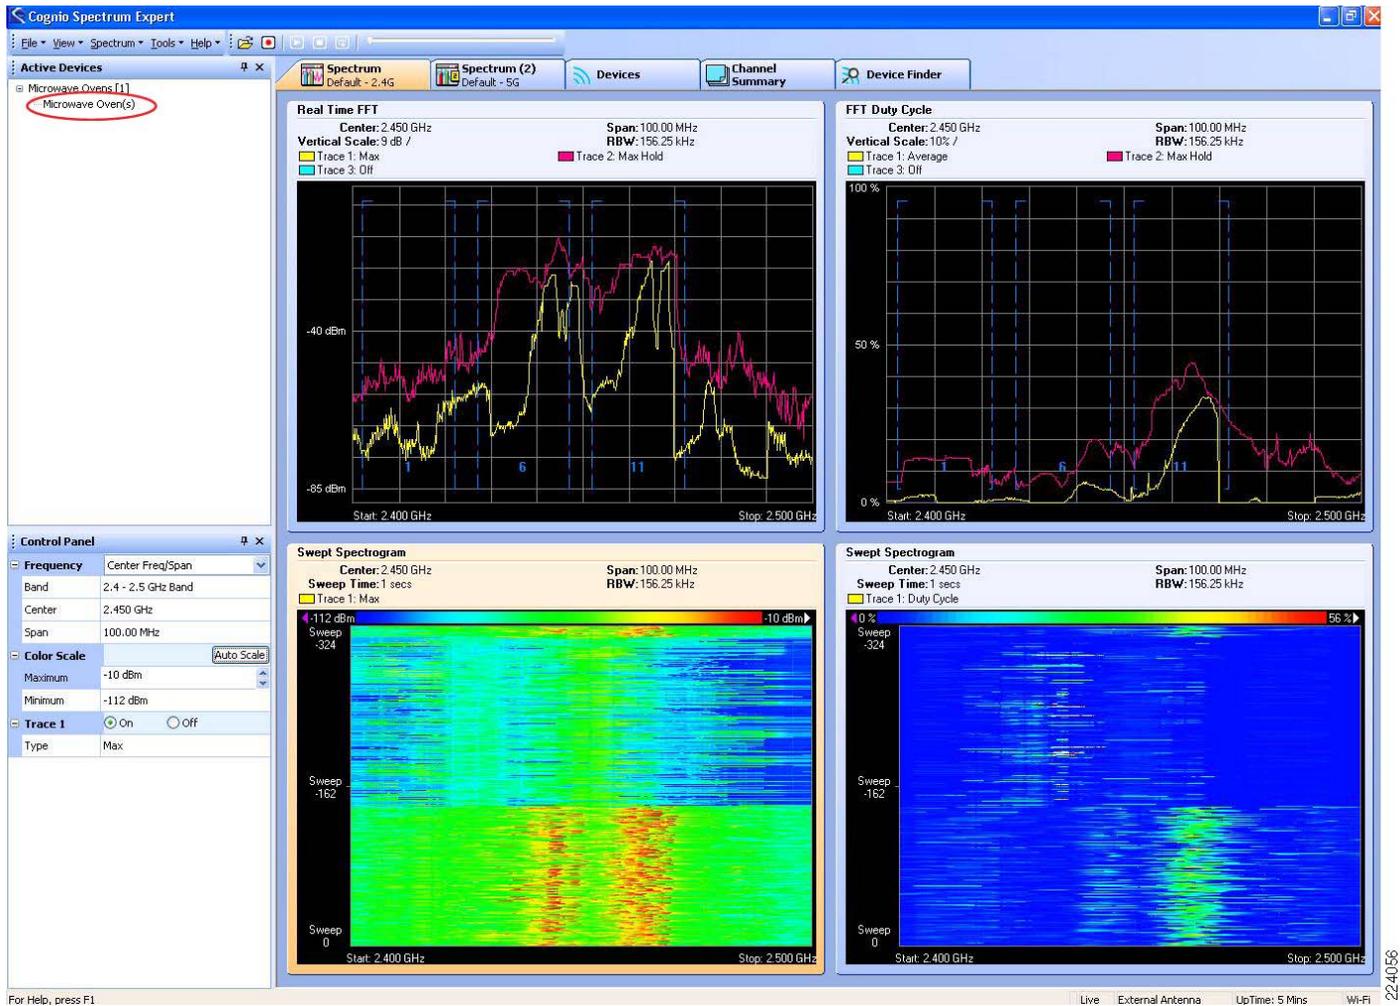


Figure 3-4 Automatically Classified Source of Interference—Microwave Oven

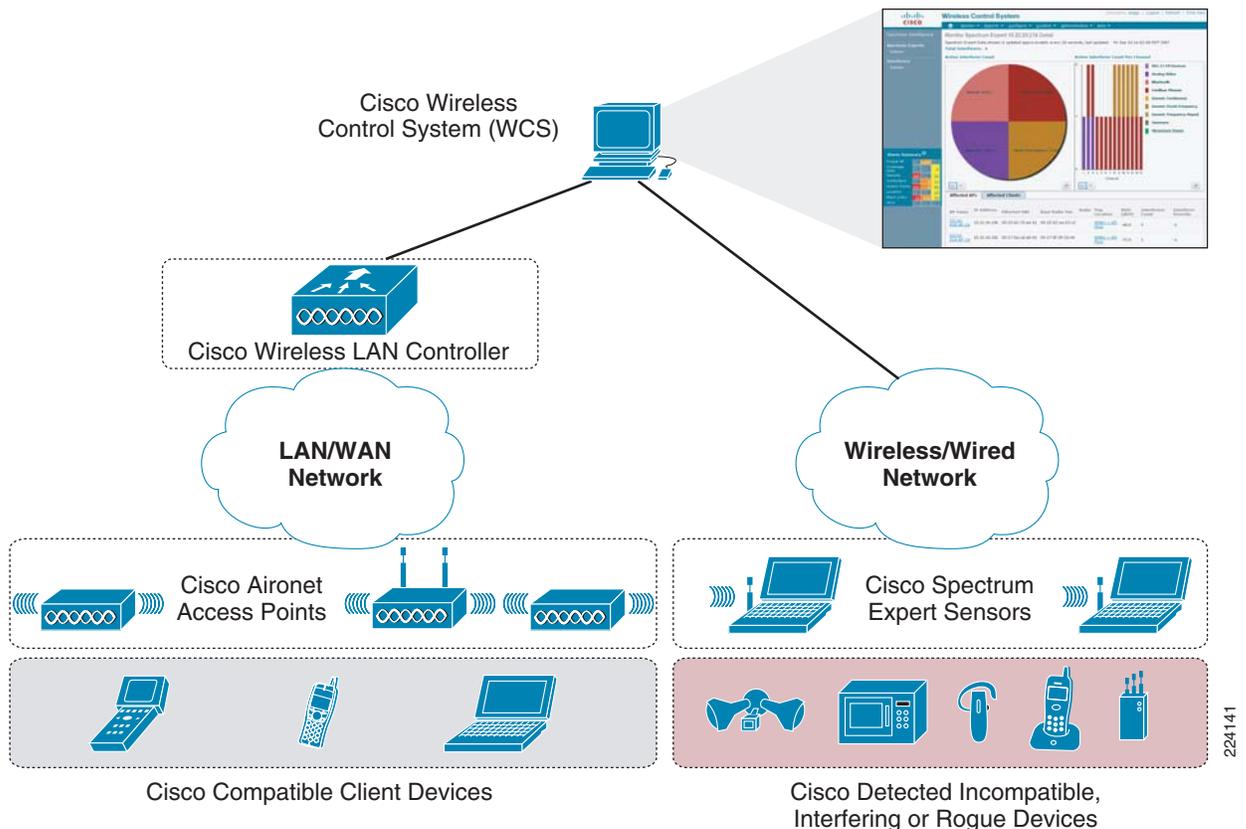



Note

A microwave oven typically does not affect Channel 1, but only affects channels 6 and 11.

In addition to using the Cisco Spectrum Expert tool directly, there is another source for obtaining current and historic information about interference. Each access point informs the Wireless LAN Controller (WLC) about interference detected every 2 to 3 minutes. This information can be consulted to provide both real time and historical tracking of such events. In multiple WLC environments as is common in most healthcare environments, the same information is available in the Wireless Control System (WCS) as shown in [Figure 3-5](#).

Figure 3-5 Information Available in the WCS



When the Spectrum Analysis sensors are not available, the APs send Radio Resource Management updates to the WLC/WCS that are analyzed for interference, and dynamic action taken. Historical reports can be produced that highlight all past interference, power, and channel changes seen on certain access points, floors, and buildings.

Effects of Radar in 5GHz Band

Figure 3-3 and Figure 3-4 show the Cisco Spectrum Expert tool monitoring the 2.4GHz band. There are reasons to use the tool to also monitor the 5GHz band. There are 5GHz wireless phones and surveillance cameras that use the same frequencies as 802.11a access points. They would be a source of interference, as would radar stations that may use 5GHz.

Radar in 5GHz must be monitored by the access points that are configured to any of the UNII-2 channels, including the channels from frequencies 5260 to 5700 or the channel numbers from 52 to 144. The reason the access point must monitor those channels is because of radar for weather and aviation. If radar is detected by the access point or other access points connected to the controllers, then the interfering access point must leave that channel in 100ms. That would force the clients to roam. Such a roam could take a longer than normal because the client may not have a current list of candidate access points to which to roam. The recommended way to handle this issue is to monitor, for a period of a week or two, the UNII-2 band. Detection of radar is logged on the Wireless LAN Controller and is reported to WCS. Radar would also be detected by the spectrum tools and could be reported to WCS. Once it is known what channels would be subject to interference and what channels see radar, then those channels can be excluded from RRM. By excluding channels with known interference, there should be fewer roams by clients and fewer channel changes by RRM or radar detection.

Another use for DCA channel exclusion is to keep RRM from enabling channels that are not supported by your clients. There are older 802.11a clients that do not support channels above 64. There are also some recent wireless clients that do not support channels 100 to 140. See [Figure 3-6](#).

Figure 3-6 DCA Channel Exclusion

The screenshot shows the Cisco Wireless configuration page for 802.11a channels. In the left-hand navigation menu, the 'DCA' option under the 'RRM' section is circled in red. The main content area displays the '802.11a Channels' configuration. At the top, there is a field for 'DCA Channels' with a list of excluded channels: 36, 40, 44, 48, 52, 56, 60, 64, 100, 104, 108, 112, 116, 132. Below this, there is a table with columns 'Select' and 'Channel'. The table shows the following channels selected (checked):

Select	Channel
<input checked="" type="checkbox"/>	36
<input checked="" type="checkbox"/>	40
<input checked="" type="checkbox"/>	44
<input checked="" type="checkbox"/>	48
<input checked="" type="checkbox"/>	52

All that is required to remove a channel from RRM configuration is to un-check the channels, then apply the update.

Effects of RF on Medical Devices

Verification that a medical device is immune to outside RF transmitters is a task that is typically performed by the bio-medical engineering team within the hospital. In most cases, the ANSI C63-18 publication describes a rudimentary process to determine if a T-PED (Transmitting Portable Electronic Device) may cause interference with a medical device under test. The document is named “American National Standard Recommended Practice for an On-Site, Ad Hoc Test Method for Estimating Electromagnetic Immunity of Medical Devices to Radiated Radio-Frequency Emissions from RF Transmitters” and can be found on the ANSI Web site for purchase:

<http://webstore.ansi.org/RecordDetail.aspx?sku=C63.18-1997>

The document describes a recommended test environment and approach for introducing a portable RF transmitter to a medical device. The ad hoc test does not offer a guarantee against interference risk, but assists in the identification of medical devices that are particularly sensitive to the RF signals generated by T-PEDs. An excellent reference for RF interference testing can be found on Cisco’s Web site at:

http://www.cisco.com/warp/public/cc/pd/witc/ao1200ap/prodlit/wrlan_wp.pdf

The healthcare organization may wish to contract with an EMC/EMI/RFI consultant, such as Kimmel Gerke Associates. Such consulting companies offer training courses for EMC in Medical Electronics which can provide the training necessary to perform basic self test and validation against medical devices.

HIPAA Considerations

Good corporate governance depends on the effective management of internal controls and on the availability, confidentiality, and integrity of information within the organization. Corporate reputation, brand preservation, and financial results all depend on the defense of business processes and on compliance with a growing array of legislation and regulation. For healthcare organizations, this includes the Health Insurance Portability and Accountability Act (HIPAA). The network has a fundamentally important role to play in HIPAA compliance, because it touches every aspect of the extended organization and connects business processes. The old, perimeter-based network security model is inadequate for managing security risks related to healthcare information. Healthcare organizations need an end-to-end, system-based approach that is integrated, collaborative, and adaptive, one that helps them better manage their network security risk while helping them meet HIPAA requirements.

In a compliance environment that contains overlapping, inconsistent, sometimes untested and often contradictory laws and regulations, organizations must increasingly turn to best-practice solutions that combat their real-world information threats while helping them meet regulatory requirements. ISO 17799 is one such framework. The Cisco Self-Defending Network provides the first line of corporate defense, because it is the foundation for the organization's data, applications, and business processes—the protection of which is a prerequisite for HIPAA compliance.

The 1996 HIPAA (which took effect in 2003) is a set of federal standards that requires healthcare organizations to implement security standards that protect (and keep up to date) patient data and to standardize on electronic data interchange (EDI). HIPAA was originally designed to speed the processing of medical claims by implementing certain standards for transmitting medical data. This of course raised information security concerns, so provisions were also made to protect the confidentiality of personal health information while in transit and while being stored. The Administrative Simplification Provisions set out the specific rules that institutions must implement to comply with HIPAA; these include the rules for EDI, for electronic signatures, and for privacy standards. Although these provisions are technology-independent, any system of information security controls that a healthcare organization implements need to be integrated and comprehensive. ISO 17799 provides an independent, internationally recognized best-practice framework for achieving these objectives, and the Cisco Self-Defending Network aligns itself with the controls recommended by ISO 17799.

Wireless Medical Devices

EMR Systems

Clinical EMR systems are at the core of driving many if not all clinical workflows within a healthcare organization. Typically, however, certain workflows lend themselves to being extended to the mobile caregiver. The most common examples are:

- Computerized Physician Order Entry (CPOE)—Allows physicians to enter prescriptions and request tests at the point of care.
- Medication administration—Reduces medication errors
- Electronic medical record—Clinical systems
- Asset tracking—Biomedical engineering

The CPOE applications typically run on a small handheld device, often a PDA. The CPOE application allows a physician or caregiver to quickly enter orders for the patients under their care. The wide range of devices supported by the clinical application provider gives a rather large variation in RF capabilities. With the advent of consumer class Wi-Fi PDAs and dual-mode phones, healthcare organizations are faced with an increasingly wide variation of wireless devices being used. Many of these devices are focused on a consumer using them in public Wi-Fi hot spots. Wi-Fi hot spots by design have no authentication or encryption and little need for fast and reliable roaming, as the end user is typically stationary. The result of using some of these consumer-class devices in a healthcare enterprise often results in poor overall performance.

A careful set of wireless standards must be developed by each healthcare organization to safeguard and set minimum standards for overall wireless security and performance. Many clinical ISVs (Independent software vendors) sponsor device fairs which allow the clinical staff to evaluate a range of IS- supported PDAs or mobile computing devices. The devices selected by the clinical staff may not meet the technical and security requirements set forth by the hospital's network IT department. It is therefore recommended to create a wireless adherence policy that outlines the minimum base requirements that the wireless device must exhibit. Without such a minimum requirements list, consumer class devices may be presented to the clinical users which are not capable of adhering to the hospitals IT security policy for wireless. Some examples of a minimum standard include:

- Authentication using EAP-FAST, EAP-TLS, or EAP-TTLS
- Encryption using either WPA or WPAv2
- Fast roaming (especially for voice) support
- WMM (Wi-Fi Multi Media) and U-APSD (Unscheduled Automatic Power Save Delivery) support
- Dynamic Power Control (reduces co-channel interference induced by clients by dynamically adjusting output power)

Clinical systems such as HIS (Hospital Information Systems) systems are also commonly accessed via a wireless device provided that the devices display characteristics are adequate to provide the full screen capabilities. To facilitate the use of HIS systems at the point of care, Computers on Wheels (COWs) are commonly used on the patient floor. As clinicians often switch from one COW to another throughout the day, IT organizations sometimes implement a Citrix-based access method. The use of Citrix enables a virtual desktop to be dynamically moved from one wireless device to another, eliminating the need to re-navigate through the various screens to reach the functions they most frequently use.

Citrix has historically exhibited a little tolerance for momentary connectivity disruptions. These disruptions are often caused by poorly-designed wireless networks where cell overlap is not sufficient or the method of authentication is not providing rapid re-authentication. The use of EAP-FAST along with secure fast roaming technique such as Proactive Key Caching (PKC) is recommended. PKC is an extension to the 802.11i standard and precursor to the 802.11r standard that facilitates secure roaming using AES encryption and RADIUS authentication. Without the use of Citrix, the bandwidth requirements for HIS systems will vary greatly depending on the clinical function or pathway used by the caregiver. When Citrix is used however, the bandwidth and latency requirements are both reduced and normalized so that traffic spikes are minimized.

The clinical software vendor should be able to provide bandwidth consumption numbers on a workflow-by-workflow basis. Using these numbers, along with a predicted concurrent use on a given access point, yields a result that can provide a reasonable bandwidth usage model. All of the application usage models should be combined to provide a reasonable usage estimate for the wireless network. Keeping track of the applications being used on the wireless network along with their traffic characteristics is recommended. This extends to version upgrades to each application as bandwidth usage may change dramatically between releases. During the pilot phase, the exact wireless devices, along with shipping operating system and device drivers, should be tested exactly as they will be used in the live network. Keep in mind that some consumer devices, while supporting the authentication and

encryption mechanisms adopted by a healthcare organization, such as WPA or WPA2, fall short in other areas. These areas may include the inability to roam between access points to varying degrees. These devices **must** be tested in the same physical environment in which they will ultimately be used.

Smart Infusion Pumps

Infusion pumps are typically thought of as life critical devices, as in many cases the accurate delivery of medication is indeed life critical. Most, if not all infusion pump vendors today, have architected their infusion pump products to function in protected mode when network connectivity has been disrupted. During the network disruption, manual protocols are used to administer new medication or changes to the infusion rate via the infusion pump. During this time, all drug advisories continue to be enforced on the pump since the last medication formulary download. Switching to a manual protocol process can be cumbersome for experienced care providers and should be avoided if possible. In the case of a locum care provider however, the manual process may simply not be known to this set of care providers. The end result is delayed delivery of medication because of the disruption of the automation provided by the infusion pump.

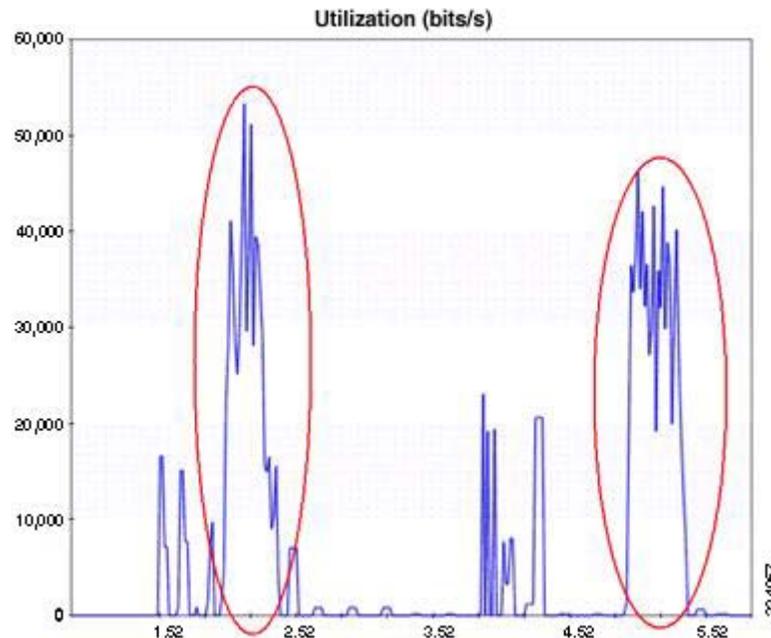
Healthcare organizations are faced with new challenges since the introduction of smart infusion pumps. To reduce medication errors, the smart infusion pump has proven highly effective in reducing medication errors through the use of drug protocol libraries. These libraries typically consist of a partial list of the overall hospital's pharmaceutical formulary which is specific to infusible medications. To keep these formularies updated on the infusion pumps throughout the hospital, a scheduled "drug protocol download" of an approximately 1 MB file is performed over the wireless network. Most infusion pump device manufacturers utilize TCP/IP for the actual download, and it is sent as a unicast stream to each pump over the course of a few hours.

From a sheer traffic flow perspective, the traffic generated is low in volume. It is however critical in some cases that the refreshes occur in a relatively timely manner. As such, careful planning with regard to download frequency and the number of concurrent downloads is essential. The frequency of these downloads is dependant on many factors and cannot typically be controlled by the network administrator. Some examples that would require a download are new drugs being introduced for treatment or a revised medication guideline as determined by the hospital formulary committee, pharmaceutical manufacturer, or FDA.

The library downloads consume relatively little bandwidth when downloading to a single infusion pump. If however all of the pumps are scheduled for a simultaneous update, the effects on other wireless applications could be substantial. Other medical devices that do not QoS mark their traffic and lack the ability to detect and retransmit dropped traffic may be at risk.

A good example of such a device would be a patient monitoring solution from a number of different vendors. These devices historically do not QoS mark their traffic, do not support WMM, and cannot retransmit traffic that may have been dropped due to interference or congestion. This issue exists even if the devices are located in a different SSID as interference or contention in the RF channel is not solved by separating devices among a number of SSIDs. The recommendation therefore is to limit the number of library protocol downloads to a handful of infusion pumps at any one time. [Figure 3-7](#) shows the typical bandwidth as measured during a formulary download to a single infusion pump. Notice the relatively small burst of traffic over a given time period does not exceed 50kbps.

Figure 3-7 Sample Bandwidth as Measured During a Formulary Download to a Single Infusion Pump



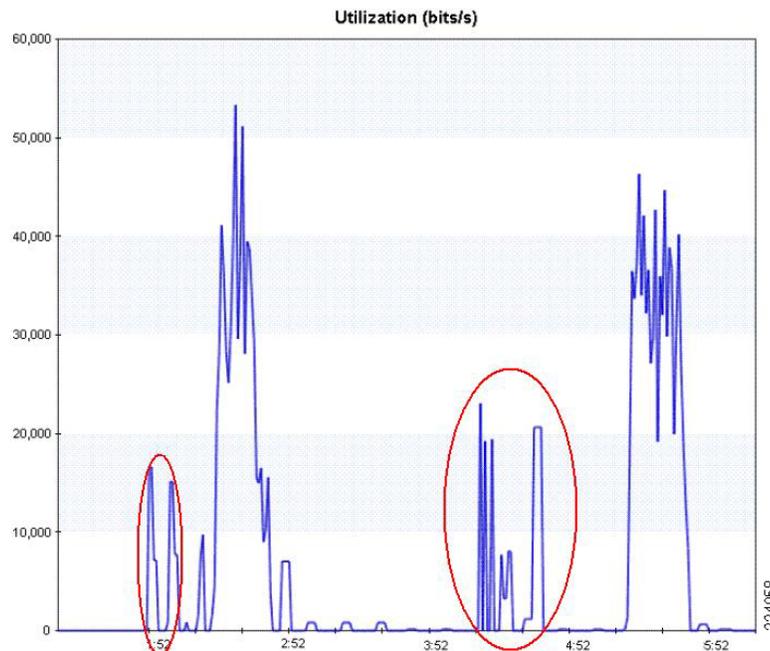
As with any wireless client, the QoS marking of the traffic originating from the wireless device is critical. In an LWAPP environment, marking the traffic as it exits the Wireless LAN Controller does little to expedite the traffic traversing the tunnel. Assigning the devices to an SSID that has platinum voice support does not mark the traffic being generated by the wireless device, nor does it mark the traffic as it enters the LWAPP tunnel. This is simply a policy assigned to an SSID that is applied to traffic matching a specified QoS marking; if not marked, then no special treatment for the traffic originated on the wireless device can be taken. It is often incorrectly assumed that simply configuring an SSID with a QoS policy automatically marks all traffic from any device using this SSID; this is simply not true. Until medical device manufacturers (MDMs) begin to properly adopt a QoS model for their critical clinical traffic, there is nothing that can be done to prioritize such traffic on the RF side of the link. The traffic must therefore contend with other less critical traffic for access to the RF link, making it effectively best effort.

Another mechanism that can be used to reduce the side affects described above is accurate RF engineering during the site survey process. Many infusion pumps still rely on 802.11b modulation techniques and therefore are limited to 11 and 5.5 Mbps using CCK modulation, 2 using DQPSK, and 1Mbps using DBPSK. As such, infusion pumps which are located within a patient room may in some cases be operating at the lowest speed within the band (1Mbps). As such, all wireless devices being serviced by the associated AP are subject to lower throughput during the time that these devices are using the RF channel.

Engineering the wireless network such that 802.11b devices achieve at least 5.5 Mbps or better, in all areas and for all medical device types, is critical. Smart infusion pumps generate small data usage spikes during drug protocol library updates which are downloaded to the pump from a vendor-specific backend monitoring system. Other data related to the performance of the pump, and the exceptions to the medication delivery protocol for the given hospital policy, is also generated and sent to a backend smart pump vendor server. This data is typically very low in quantity and does not pose a problem to a properly engineered wireless network. Delivery information as to the time, quantity, and caregiver delivering the information, however, must reach the backend system on a timely basis. This is necessary as many

healthcare organizations feed this information to their clinical system to update the patient's EMR. Figure 3-8 shows the traffic typically generated when a clinician initiates the delivery of medication, as well as the information collected by the pump and fed to the EMR system via the pump's backend server.

Figure 3-8 Traffic Typically Generated When a Clinician Initiates the Delivery of Medication



Tracking the location of smart infusion pumps by the bio-medical engineering department is often an important consideration. The device can be pinpointed provided that the device is powered on and history has been enabled within the location tracking system in use. The Cisco 2700 location appliance can provide historical information as to its last location when powered on. If, however, the device was powered off, moved, and remains powered off, accurate location information is not possible through the use of the onboard 802.11 wireless NIC. If, however, the device was tagged with an active RFID tag, such as Pango or Aeroscout, the location can be tracked without a dependency on the device being powered on.

Roaming with infusion pumps is not generally an issue as the pump is often stationary when compared to a caregiver. However, if the pump becomes mobile due to a patient leaving their room or being transported to an ancillary department for a procedure, roaming may cause momentary interruptions in connectivity. The performance of roaming is vendor (Medical Device Manufacturer) dependant, and usually revolves around issues with re-authentication or the wireless client's driver and its roaming algorithm.

In 802.11-based networks, the decision to roam to another access point is the responsibility of the wireless client device. If the roaming algorithm is "stickier," the pump may not roam until it is too late and so disassociates from the access point. Then when the roam does occur, a slow channel scan or non-optimized authentication process may begin. The bottom line is a slow roam or possibly momentary disruption of network connectivity. Again, due to the nature of most infusion pumps on the market today, the pump operates in a protected mode and continues to deliver the medication at the prescribed rate. Information collected is buffered and the performance metrics uploading to the EMR at a later time.

Many infusion pump vendors request that their pumps be placed on an isolated SSID, one that is only shared by like vendor pumps. This practice stems from the fact that some medical devices, specifically patient monitoring systems, generate a large amount of broadcast or multicast traffic. Because the IP

stack in the infusion pump is constrained due to limited memory and CPU capabilities on the imbedded device, reducing excess broadcast through SSID isolation is a common solution for preventing buffer overflows and increasing battery life.

From an RF channel perspective however, a broadcast generated on any SSID is still seen at an RF or PHY level on each wireless NIC associated to the same AP regardless of the SSID used. The wireless NIC decides if the broadcast needs to be forwarded up the stack for processing. Depending on the design, this may or may not induce additional CPU and buffer overload concerns. It is therefore recommended to request from the infusion pump vendor the list of requirements as it relates to SSID isolation as a method of controlling broadcast traffic. An 802.11 wireless network has a certain amount of management traffic that is typically transmitted by the access-point. Standards-based 802.11 management frames provides a number of capabilities to the overall network.

Many wireless vendors and other wireless devices may generate management traffic that does not conform to the 802.11 standard. In some embedded systems, the parsing of these management frames have been know to cause the wireless NIC to hang. The cause of the problem is often placed on the devices that generated the 802.11 management traffic, but in actuality the culprit is a poorly designed parsing algorithm on the embedded system. There may be times where an 802.11 management frame becomes corrupted due to outside RF interference. If the parsing algorithm is not robust enough to ignore those frames that it is not programmed to understand, then a wireless NIC driver hang may result. If this does occur, its important to first contact the device vendor of the affected system, and secondly re-examine any changes that may have been implemented on the wireless network.

Signal level recommendations vary from medical device to medical device. This is primarily due to the fact that the receiver sensitivity may vary widely from vendor to vendor and even vary somewhat between production runs of the wireless components themselves. Some pump vendors recommend that the RF network be engineered so that each pump receives -70dBm of signal and a minimum of 25dB Signal to Noise Ratio (SNR) in all areas where the pumps operate. One often overlooked items is that the site survey may have been done with a laptop or other device where the antenna is located in an optimum location (external or in the LCD display). Medical devices on the other hand may not have taken the optimum placement of the antenna into consideration during the product design phase. The end result, from the perspective of the smart pump, may be a signal level that is lower then the recommended value even though the site survey results are adequate. Always allow for variations in the wireless capabilities of a medical device. In summary:

- Engineer network for 25dB and -70dBm in all areas where infusion pumps operate.
- Limit the simultaneous number of infusion pumps that receive formulary updates to 5 or less per AP.
- Drug formulary downloads on a pump-by-pump basis are typically very small in size, relatively speaking. Other devices that may be affected are most commonly devices that do not implement WMM or QoS and may have limited or non-existent retransmission capabilities.
- Isolated or dedicated SSIDs for infusion pumps may be recommended by some vendors to protect limited CPU and buffer resources in the infusion pump.
- Infusion pumps are indeed life critical, but in most vendor implementations operate in protected mode and are therefore able to infuse medication without network connectivity.
- Roaming of infusion pumps is generally not an issue, but is dependant on the roaming algorithm implemented in the wireless driver in the infusion pump.

Patient Monitoring Systems

Patient monitoring systems are designed to monitor the physiological condition of a patient using a number of different metrics. Many of the higher-end patient monitor devices can operate in an autonomous mode and by themselves detect abnormalities in the patient being monitored. Others may rely on the computational capabilities of a backend system which analyses the telemetry and alerts when one of the monitored waveforms is outside of the normal bounds of acceptable limits.

Understanding the history and evolution of patient monitors is critical to a better understanding of their current technological state and hence the rigid network requirements that they present to a wireless network. Historically, these devices were only available in areas where significantly injured or ill patients resided. The emergency department, operating room, or intensive care units are a few of the most common areas.

The need for the caregiver to monitor any number of patients under their care required that the patient monitors be connected to a central station. This connectivity did not begin with Ethernet, nor was the communication IP based. Instead, these devices required dedicated cabling between the patient monitor and the central station. Typically, the communications protocol was a serial rs-232 link which utilized a vendor-specific serial multiplexing device which combined many of the feeds and was capable of displaying multiple patients on a single station monitor.

This private network provided the patient monitor vendor with a high degree of assurance that the waveforms collected at the point of care would in fact be presented without disruption or a delay of even a few milliseconds. Due to the high cost of these proprietary, serial-based solutions, combined with the relatively low cost and pervasiveness of Ethernet technology, eventually all vendors implemented Ethernet into their products.

Because of this dedicated serial network approach, many patient monitor vendors required that the Ethernet based architecture of this network also remain separate. In fact, during this time, Ethernet networks throughout the hospital were not very pervasive. It would not be unreasonable to assume that for some hospitals, the first Ethernet networks installed were the isolated patient monitoring Ethernet networks. This departmental approach to putting a network where needed by installing an Ethernet hub in the ceiling or under a desk provided all that was needed at the time.

As patient monitoring vendors enjoyed the isolation that these private Ethernet networks provided, the rest of the healthcare organization was becoming wired. With the advent of clinical systems and the Electronic Medical Record (EMR), data connectivity between these historically separate networks was required. To provide a bridge between the patient monitoring network and that of the EMR system, a dual-homed gateway was commonly used to forward telemetry data and alarms to the clinical system. This provided the needed connectivity, but allowed the IP isolation of the patient monitoring network to remain in place.

During all this time, the need to detect a dropped packet, or retransmit data that was lost due to congestion or other network issues, was never added to many of the patient monitoring product lines. To further complicate the matter, multicast was used by each patient monitoring device to send the waveforms not only to the central station, but also to any other patient monitor that requested remote viewing of a particular patient monitor within the department. This functionality was needed to provide the caregiver with immediate access to a waveform of another bed during an alarm event.

Let us fast forward to today. Healthcare organizations worldwide now have Ethernet networks that are pervasive and provide a high level of availability due to various levels of redundancy and end-to-end QoS. These organizations do not want to incur the added costs of parallel standalone networks and expect that the patient monitoring solution utilize the existing infrastructure. From a patient monitoring vendor standpoint, the availability of such networks can vary widely between healthcare organizations. Their goal is to deliver a consistent solution that meets the same level of service that has been historically provided by their product set.

With the advent of 802.11 wireless networking, a migration similar to that seen in the early days is once again occurring with wireless. The wireless networks this time are not meant to replace the wired Ethernet network, but rather to augment it and provide consistent connectivity when the patient is mobile. Often a critically-ill patient needs the services of an ancillary department such as radiology. Many of the patient monitoring systems available today have the ability to become mobile and dynamically switch to an 802.11 wireless network. The device is sometimes connected to the bed, or in the case of some of the smaller devices, put on the end of the bed while the patient is moved throughout the hospital.

During this time, the same wired line equivalence is expected of the wireless network that existed with that of the isolated patient monitoring Ethernet network. The challenges arise when such assumptions are made that the wireless network has the same capabilities that once existed in the private network. The simple ability to detect frame loss and retransmit during times of RF interference is simply not part of the firmware in many patient monitoring products today. Add to this the fact that the 2.4GHz spectrum has a high noise floor and may in fact be shared with other applications and one can quickly conclude that there are significant challenges in providing wireless connectivity similar to that of a switched Ethernet network.

As discussed in [Chapter 2, “RF Design Considerations,”](#) the FCC states that any device using unlicensed frequencies must tolerate any interference caused by other devices operating in the same unlicensed frequency. Interference in any RF environment should be expected and therefore planned for. Designing a wireless network that provides connectivity similar to that of an isolated wired network may prove difficult, but there are some approaches to such a design that should be considered.

First, from an interference standpoint, there is little that can be done to prevent interference, so the next best approach is to utilize an 802.11 band that is less prone to such interference. Selecting 802.11a provides the patient monitoring system with a wider set of channels than that of 802.11b/g, and at the same time typically has a lower level of noise. A good choice is the AP1250 series of access points because the AP1250 series supports all 21 channels of 802.11a and therefore offers the most flexibility in the design of the wireless network. It is however important to verify that the patient monitoring vendor supports both 802.11a and the full set of 21 channels. If it supports 802.11a, but not the full set of channels, the only recourse is to disable those channels on the wireless network to ensure that a patient monitoring device receives consistent wireless connectivity while roaming throughout the hospital.

Roaming is another concern. As discussed in [Chapter 2, “RF Design Considerations,”](#) the decision to roam to another access point is controlled by the firmware running on the wireless card designed into the patient monitoring product. Some roaming algorithms are “stickier” and do not decide to change to another access point until its too late, usually resulting in some data loss. Because many patient monitoring devices have no way to detect and retransmit the telemetry data, gaps in the waveforms may appear on the central monitoring station. To optimize this, it is recommended to create smaller cells (802.11a if possible) thus providing the patient monitor devices with a denser choice of access-points from which to choose.

The use of multicast can also be challenging in that the 802.11 specification requires that broadcasts and multicasts be sent at the lowest configured speed. This requirement stems from the fact that a broadcast (and therefore multicast) frame is intended to be received by all hosts within the broadcast domain. The only way to help ensure that this is possible is to transmit it at the lowest speed in the hope that hosts that are far away from the access point or transmitting host receive the broadcast frame. Therefore it's important that the lower transmit speeds, such as 1Mbps and 2Mbps, be disabled on the access points throughout the facility.

Disabling the lower speeds also has a performance advantage to other hosts within the cell due to the decrease in retransmissions caused by wireless clients in other areas whose signal is generating interference.

There are a number of different traffic flows to consider in a patient monitoring network design. The most common, however, is the patient device to the central station. Depending on the vendor, each patient monitoring device transmits between 2-6 (or sometimes more) packets per second. The packet size also varies from vendor to vendor, but typically is in the range of 200-400 bytes. Using these theoretical numbers from our example, each patient monitor generates raw data of upwards of $(6 * (400 * 8))$ bps or in our example, 19,200bps. Added to this is the frame overhead of 34 bytes for the 802.11 MAC header, 20 bytes of IP header, and 8 bytes of UDP header, totaling 62 bytes or 496 bits. This brings the grand total up to just under 20kbps per patient monitor. Your results will vary depending on the vendor and the number of leads (waveforms) being monitored.

While this number seems quite low, and it is when viewed by itself, the aggregate bandwidth of a 100 beds, each generating 20kbps of multicast, brings the total to just under 2Mbps of combined multicast traffic.

From the perspective of multicast, it's commonly understood that any hosts that has requested the multicast feed should receive it by requesting access to the Multicast Group using the IGMP protocol. With some patient monitoring vendors, this is not the case. Going back to our patient monitoring history lesson described earlier, the fact that the network was isolated allowed the vendors to partially implement multicast. What this means is that some vendors do not support dynamic group membership via the IGMP join/leave processing. Since the patient monitoring devices evolved on a standalone network, there was no need to add such complex logic to the product set as all multicast traffic was seen by all other patient monitoring devices connected to the isolated network.

To implement similar functionality on a wireless network, all multicast traffic must therefore be seen by all other patient monitoring devices, regardless of their location within the hospital. Using our example above, this results in the fact that all APs supporting such vendor implementations must continuously stream the 2Mbps of patient monitoring data, regardless of whether or not there is a patient monitoring device viewing the results of a particular multicast stream. It is therefore important to allocate this overhead into the bandwidth requirements of the overall wireless network.

It is disappointing that to date little attention has been given to implementing QoS in the patient monitoring devices themselves. With data that is indeed life critical, the absence of WMM and 802.11e support is indeed troubling for a biomedical network engineer. Over time however, it is expected that the industry will respond and implement 802.11e and/or WMM along with IGMP multicast support natively within their products.

Providing QoS marking or traffic shaping for traffic in the downstream direction from the perspective of the AP (AP⇒Patient Monitor) does little to provide assurance for traffic generated at the patient bedside (Patient Monitor⇒AP). The best advice is to provide a dense AP infrastructure that reduces the likelihood of RF contention with other wireless clients.

From a security perspective, the authentication and encryption of the patient monitoring devices vary widely from vendor to vendor. At a bare minimum most if not all vendors support static WEP. Other vendors support WPA-PSK and various forms of EAP. Many however do not support the various techniques for fast secure roaming. As a result, the method selected for encryption and authentication directly affects roaming results, and can be exaggerated in certain areas such as Elevators where the roam may be across wireless controllers.

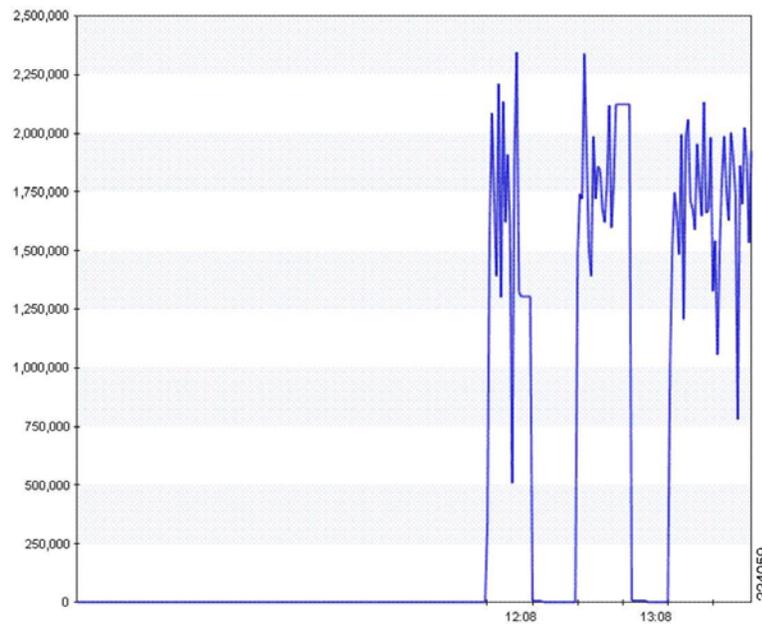
Portable Radiology and Cardiology Devices

The use of 802.11 wireless connectivity for portable diagnostic medical equipment has become commonplace in many hospitals. The advantage is to allow various studies to be uploaded to a backend PACS or cardiology systems for archival and diagnostic purposes. This improves the workflow of the mobile clinician, allowing the studies acquired in the patient room to be offloaded from the device as the technologist is executing their examination worklist. In the case of portable digital X-rays, the images

can exceed 10MB and a single exam and commonly comprises a number of separate images. The result is a study that in many cases can exceed 100MB. Transferring this to a backend PACS system over a wireless network while the technologist is roaming both horizontally and vertically within the building can sometimes be challenging.

From a pure traffic perspective, DICOM-based traffic is typically treated as best effort. The studies must therefore contend with other wireless traffic generated in a given 802.11 service area. This is because many portable device vendors do not mark their traffic as mission critical data and hence remain unmarked from a QoS perspective. As Figure 3-9 shows, traffic spikes in excess of 2.5Mbit/s can result as the acquired study is transmitted to the backend PACS system.

Figure 3-9 Portable Digital X-Ray Wireless Bandwidth Usage—Spikes in Excess of 2.5Mbit/s



As stated, since most medical devices today transmit their data as best effort, it is not unreasonable to assume that in a poorly designed wireless network, such an upload could cause network congestion and negatively affect other applications which may be critical to patient care. In general, uploads should be treated identically to a batch file transfer and therefore receive best effort service. The TCP/IP rate limiting characteristics of TCP slow start and congestion avoidance mechanisms which are native to TCP/IP effectively rate limit the transfers, as well as allowing the transfer to restart during momentary disruptions in connectivity. It is important to note that some more critical wireless applications, such as patient monitors which do not support WMM or QoS may be contending for network resources. This serves as a reminder that a medical grade network requires advanced planning and consideration.

One approach is to ensure that the wireless network design has no area that provides less than 11Mbps, and if possible provides even higher data rates. This is based on the premise that it is better to get the upload completed as quickly as possible, therefore reducing contention for the RF channel. Locking the client to a lower speed, such as 1Mbps, is not recommended as during the time that each frame is transmitted at 1 Mbps, all other clients do not have access to the RF media or channel. This is equivalent to a CB radio operator talking at a slow speed and preventing others from using the channel to communicate. Ensuring that all wireless service areas support faster speeds eliminates the bottleneck of the slower talkers.

Another approach is to determine the band where critical applications reside. Moving the less critical wireless clients to the other 802.11 band would provide the necessary isolation. The bottom line is that some life critical applications do not provide any means to categorize their traffic as critical. The end result is that both critical and non-critical wireless clients must share the RF medium and are therefore treated equally.

Medical Carts—COWs

Medical carts, often referred to as computer on wheels (COWs), have become more and more prevalent in hospitals as wireless has become pervasive. A COW is a full size computer, tablet PC, or laptop that is attached to a mobile cart and powered by a rechargeable battery (in some cases a UPS). These carts may also have attached medical instruments. For network connectivity, these COWs may plug into Ethernet wall jacks, but for the greatest flexibility, they are more commonly 802.11 enabled. COWs have several key characteristics:

- Size of the cart—The cart must be able to move into the small work spaces that require the use of the COW.
- Ease of use—The COWs need convenient and user-friendly instruments, power, plugs, and adjustments.
- Work space—The cart should provide enough work space and storage for daily use.
- Mobility—Mobility refers to the ease of movement of the cart, not wireless mobility.
- Ergonomics—The cart should be adjustable for different users and types of use, including all equipment on the cart such as the keyboard, monitor, and other instruments.
- Battery—This is a key feature that should support several hours of use, but also support easy recharging.
- Safety features—The carts need to be sturdy, lockable into a stationary position, and minimize any threat of tipping since they have to be adjustable to suit the varying heights of the users.

The true benefit of the COW is to empower the caregiver with access to patient information at the bedside or point of care. This accessibility translates into real-time access to information to support critical decision making, reduces medical errors, and eliminates redundancies in current hospital workflows. Paper documents can also be drastically reduced. Time savings, work efficiency, improved care, and the ability to create a more accurate patient record are the main benefits of being able to pull information from multiple sources without digging through paper records, especially because the information is presented in a logical fashion. The usability features of the cart are critical to avoid staff reverting to paper trails. Factors such as the mobility of the cart, having reliable power and network connectivity when needed, etc., are key elements that cannot be neglected.

Once the COWs are operational, assessing the applications used on the COW determines the network characteristics that drive the wireless network design. There can be a range of applications, such as VoIP applications and video streaming applications such as employee training or patient education. A clear understanding of these applications and the network metrics they require are key to designing a wireless network to provide adequate service to the diverse set of applications. From a QoS perspective, many of these applications may not natively mark the traffic they generate. Providing a mechanism to QoS mark the traffic being generated by the wireless COW can add significant benefit to key clinical applications. Therefore the PCs running on these COWs should be installed with Cisco Security Agent (CSA) which can support QoS markings and at the same time reduce security threats to the COW. The ability of CSA to QoS mark traffic generated can be configured at a granular level consisting not only of the TCP port, but also the destination IP address as well as the executable running on the workstation.

One last unique consideration for the COW is the antenna design and its overall effectiveness. Wireless networks must be able to reliably receive and send data in any physical location throughout the patient floor. In some situations it is possible for the AP not to be able to receive data transmissions from the wireless client even when it is within the service area of the AP. Most notebook wireless adapters use a diverse omni-directional antenna system that is optimally located in the LCD display panel. You may be able to improve the overall performance by replacing your COW's workstation wireless network adapter with a low cost external antenna. These high-gain, low-cost antennas can provide increased performance due to its external location coupled with their ability to provide diversity. Using wireless chipsets that provide Cisco CCX extension is also recommended. The use of CCX based wireless chipsets and drivers can provide an increased level of performance on the wireless network.

RFID Devices

RFID can be used to locate people, assets, and supplies. These devices have a very limited effect on overall network and application performance. The Cisco 2700 Location Appliance can track up to 1500 devices simultaneously. Each of these devices typically sends a packet with a maximum size of 256 bytes. Contained in these small packets is data such as temperature, velocity, device pressures and so on. These periodic updates are usually configurable and even dynamic in nature. If a device senses that it is in motion, the update interval can decrease for increased location accuracy; while at rest, it can send fewer updates to the network. The overall impact to the network can therefore vary and should be considered for any implementation. A typical RFID device that is used solely for tracking assets is configured to update the network every 3 to 5 minutes, with an average data packet of 30 bytes of information. Although this can change while the asset is in motion, it is more common for these assets to remain at rest in a closet or other storage area. Given this criteria, it is clear that even with 1500 active tags with a maximum update packet, the bandwidth is minimal, providing minor impact to a properly-designed wireless network.

Communication Devices—7921/Vocera/Future of Dual Mode

Communication devices are essential in a medical environment to maintain communication and receive current information. As the staff in a medical environment are highly mobile by nature, support for communication devices offered over a Cisco Unified Wireless Network solution keeps everyone connected. Communication devices can cover a broad range, including data devices such as a wireless PC or a wireless PDA for data applications as well as other forms of healthcare applications. Data applications running on portables are covered as a general category in [Medical Carts—COWs](#). Communication devices may also include legacy wireless devices such as pagers or cellular phones. The focus of this section is to further provide information on wireless devices that enable VoIP communication within the medical environment. [Table 3-2](#) provides the common VoIP communication devices that are seen in a hospital environment. It does not include 802.11-enabled cellular phones that may primarily be used for data access.

Table 3-2 VoIP Communication Devices in Hospital Environments

Type of Device	Users	Description
Cisco 7921G	Caregivers, physicians and support staff	The Cisco Unified Wireless IP Phone 7921G is an easy-to-use IEEE 802.11abg wireless IP phone that provides comprehensive voice communications. The phone supports a host of calling features and voice-quality enhancements using WMM.

Table 3-2 VoIP Communication Devices in Hospital Environments

Vocera B1000A/B2000	Caregivers, physicians and support staff	Vocera communications badge offering a lightweight, wearable, and voice-controlled device using 802.11b/g. The B2000 supports WMM.
ASCOM i75	Caregivers, physicians and support staff	Ascom VoWiFi device that is CCX v2.0 compatible and built with a Medic version to support the healthcare sector. The device is 802.11b/g and WMM compatible.
Dual Mode with Unified Mobile Communicator	Physicians	Cisco Unified Mobile Communicator is an easy-to-use software application for mobile handsets that is currently supported on models of Blackberry and Symbian OS phones. See the following link for devices supported: http://www.cisco.com/en/US/products/ps7271/products_user_guide_list.html
Laptops with Unified Personal Communicator	Clinicians including physicians	Cisco Unified Personal Communicator is a powerful communication software application in the Cisco Unified Communication family of products that enables voice, video, Web conferencing, and other key features to effectively communicate from the PC.
Laptops with communication applications such as Skype	Personal or guest access	Wireless devices used by personal users in a hospital may include a host of software applications that run on PCs that allow users to make telephone calls over the Internet. In these applications, video may also be used.
Dual mode phones with VoIP applications	Personal or guest access	Same as software applications that run on PCs, there are a host of software applications that run on wireless 802.11-enabled PDAs that make telephony calls over the Internet, such as Fring.

There are many different classifications of wireless clients as discussed previously in this chapter. Since these devices have very different traffic patterns and assumed levels of reliability and predictability, it is important to discuss outside influencers which can negatively affect the reliability of some wireless-based medical devices.

One common example is that of voice traffic. Within a healthcare environment it is critical due to the mobile nature of caregivers and their ever increasing reliance on connectivity throughout the entire healthcare facility. As such, it is critical that the VoWLAN network be architected end-to-end so that high availability can be achieved within the hospital and associated outlying buildings. To provide guidance in the proper deployment of VoWLAN solutions, Cisco has published a design guide that covers 7921G and Vocera endpoints. This Unified Communication Cisco Validated Design document as well many others can be found at <http://www.cisco.com/go/cvd>. For the Cisco 7921G specifically, Cisco has listed the 7921G with the FDA for use with what is defined as Hospital Information Systems. This includes nurse call, clinical information systems, financial reporting systems, etc. See the FDA regulations at 21 C.F.R. §§ 862.2100, 890.3725, 892.2010, 892.2020 and the FDA device listing database at <http://www.fda.gov/cdrh/reglistpage.html>.

Cisco further recommends that wireless systems be integrated in accordance with the US FDA document, Draft Guidelines for Industry and FDA Staff Radio Frequency Wireless Technology in Medical Devices (<http://www.fda.gov/cdrh/emc/>).

The use and variety of wireless-enabled devices continues to grow because of a variety of drivers, including:

- Hospital staff (caregivers, physicians, and support staff) are naturally mobile.
- WLAN technology and the expansion in hospitals is providing ubiquitous coverage in hospital environments.
- The increase in the number of dual-mode devices such that professionals will see an increasing blur between personal and work devices.
- Affordability of these devices is making it more economical for everyone to have including staff, patients, and guests.
- For the professional the convenience of having single number reachability further made possible through Cisco Unified CallConnector Mobility.

Given the large variation of devices seen in the hospital environment, some key network considerations must be addressed. Finding a common framework to support variation is a challenge as devices may and will behave differently:

- **Bandwidth**—The bandwidth for VoIP calls is typically steady state, calculated based on the codec used. Two common codecs are G.711 and G.729. Since voice samples are taken at a predefined number of milliseconds, the data rate is constant. For example G.711 sample rates are often 20ms and G.729A is typically 10ms. Another factor for the overall bandwidth is the IP header.
 - For example, G.711 has an input bandwidth of 64kbps with a sampling rate of 20ms, which translates to 50pps (packets per second) with a payload size of 160 octets plus a 40 octet fixed size overhead for IP/UPD/RTP headers. The bandwidth results in 80kbps.
 - Using G.729a, which has an input bandwidth of 8kbps and sampling at 10ms, this translates into 50/pps with a payload size of 20 octets. Adding in the IP/UDP/RTP header, the resulting bandwidth is 24kbps These per call bandwidth settings should be used in the overall bandwidth sizing effort.
 - Another important parameter in VoWLAN planning is call capacity—the number of simultaneous VoWLAN calls that can be supported in an area. This value can vary depending upon the RF environment, the VoWLAN handset features, and the WLAN system features. For example, the VoWLAN maximum capacity for a Cisco Unified IP Phone 7921G using a WLAN that provides optimized WLAN services (such as the Cisco Unified Wireless Network). See [Chapter 2, “RF Design Considerations”](#) for call capacities.
- **Quality of Service**—Due to the sensitivity of VoIP traffic to factors such as packet loss, jitter, and delay, the Cisco Unified Wireless Network suite of products supports Wi-Fi MultiMedia (WMM) standards as the framework for the QoS designs. As a range of other wireless devices will also be introduced in the hospital network to support voice, determining the compliance of these devices to the WMM standards helps to create a more stable environment to deliver voice. For example, the Cisco Unified Wireless Network supports the use of the Cisco 7921G.
 - **WMM**—A QoS system based on IEEE 802.11e that has been published by the Wi-Fi Alliance, and WMM Power Save, as well as Admission Control.
 - **Traffic Specification (TSPEC)**—An 802.11e admission control model that allows an 802.11e client to signal its traffic requirements for prioritized access to the AP through the use of two variables, the EDCF option and the controlled access options provided by the transmission opportunity (TXOP).
 - **Enhanced Distributed Channel Access (EDCA)**—A method that allows high-priority traffic a higher chance than low priority traffic to wait less before sending a packet. This model requires the TSPEC parameters defined on the endpoint. For voice traffic, this feature allows voice packets to be delivered with less delay and jitter.
 - **QoS Basic Service Set (QBSS)**—Information elements that are advertised in beacons by the AP and the parameters are used by the 7921G phone to determine the best AP to associate with.

- Unscheduled automatic power-save delivery (U-APSD)—A feature that allows the voice client to synchronize the transmission and reception of voice frames with the AP, which allows the client to enter power-save mode between the transmission/reception of each voice frame tuple. Another added benefit is to increase call capacity due to efficiencies gained by the AP.

These factors and more should be examined as the QoS model is determined to best service VoIP traffic on a hospital premise.

- Security—Another key element to evaluate for endpoint support. The key categories for support include Authentication options:
 - Lightweight Extensible Authentication Protocol (LEAP) Authentication
 - Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FAST)
 - WEP/WPA/WPA2 Shared Key
 - Encryption options
 - Wired Equivalent Privacy (WEP)
 - Temporal Key Integrity Protocol (TKIP)
 - Advanced Encryption Standard (AES)
 - Fast roaming protocol options
 - Cisco Centralized Key Management (CCKM) (CCKM is supported with TKIP/WPA only; AES/WPA2 is not supported)

Since the endpoint capabilities vary, the recommendations for security and other capabilities may also vary. The recommendation for 7921 is:

- WPA—Authentication and key management
- TKIP—Encryption
- CCKM—Fast Roaming protocol

Table 3-3 Setting Recommendations for Vocera

Authentication	Encryption	Message Integrity Check
Open	None, WEP64, or WEP128	N/A
LEAP	TKIP-Cisco, WEP64, or WEP128	N/A
WPA-PEAP (MS-CHAP v2)	TKIP-WPA	MIC
WPA-PSK	TKIP-WPA	MIC

- Coverage/Roaming—This area focuses on considerations that need to be examined to gain the necessary coverage for voice to meet the business needs of a hospital. Proper coverage for voice dictates that there is a 20 percent (2.4 Ghz) and approximately 15 – 20 percent (5 GHz) which is above and beyond WLAN data design guidelines. In addition, the optimal VoWLAN cell boundary recommendation is -67dBm and the separation of cell should be 19 dBm. For voice traffic, a unique SSID should be dedicated for VoIP traffic. Proper site survey is essential to plan for any site deploying voice. In addition to these basic factors, additional considerations include:
 - Coverage models to support elevators, stairwells, parking lots, and outside zones since a large campus may need to support staff in all these locations for reachability. There are assessment tools such as A2Q (assessment to quality) and site surveys that should be strictly followed.

- Use the tools offered in WLC to collect statistics for troubleshooting, such as AP delay, packet loss, radio statistics of the overall RF environment, as they can be helpful in finding trouble spots. WCS also provides historical reports that correlate user issues with network issues that find potential capacity issues. To provide an overall view of WLAN health, the WCS provides reports of the AP power and channel changes over time.
- Another very critical factor for voice quality is based on the amount of movement of the person talking. Devices such as laptops may need mobility, but once in use the likelihood for movement drops. However, Cisco 7921G, Vocera badges, and dual-mode phones are highly likely to experience movement. This mobile characteristic drives the need for fast and secure roaming. The client is the device that determines when to roam. The client uses factors such as data retries, RSSI, SNR, and other schemes. The network can aid with the speed in which the roaming can occur. Once the new AP is chosen, there is also a need to reauthenticate with the new AP as security must be preserved after the roam has occurred. For fast secure roaming, a method offered through the use of Cisco Centralized Key Management (CCKM) and Proactive Key Caching (PKC) allows a WLAN client to roam to a new AP and re-establish a new session key, known as Pairwise Transient Key (PTK), without having to perform a full 802.1X/EAP reauthentication with a radius server.

Bar Code Scanners

Medical Administration Check is a clinical application that is used to confirm patient identity, the correct drug, correct dose, correct time, and correct caregiver (trained to administer and monitor for side effects). This is all facilitated through the use of bar code scanners and devices specifically made for the healthcare market. Symbol, MobileComputing, and Intermic are some vendors common in this space. For standalone barcode scanners, some use 2.4GHz Frequency Hopping to tether the device to the COW running the pharmacy application.

Motorola Symbol provides handheld mobile computing devices that utilize 802.11 wireless technologies as the sole method of network connectivity. Motorola Mobile Physician Rounding Solutions allow doctors to enter patient observations directly into the EMR at the point of care. Many of these devices are handheld mobile computers that are keypad or penpad based. Besides the keypad or penpad data entry features these units are often coupled with peripherals that can read 1D & 2D BAR codes, magnetic stripes, and smart cards. Their handheld computers can also be paired with printers to provide remote WLAN printing. Be aware that the pairing may utilize Bluetooth, which may interfere with nearby WLAN devices. The Symbol client devices can be used for any number of applications, including patient records, billing systems, claim systems, maintenance systems, and logistics. The patient care applications include prescription entry, ordering medical tests, specimen tracking, prescription drug tracking, asset identification, and patient identification. The patient care applications are generally real time and mobile applications that used the Cisco Unified Wireless LAN. These applications have lower bandwidth and latency requirements than that of voice and video. But due to their critical nature, they still need to be assigned QoS-enabled SSIDs. These applications should be assigned to the best effort queue.

Many Motorola Symbol client devices are Wi-Fi certified. The handheld units that are Wi-Fi have been tested against Cisco access points, as the Cisco access points are part of the Wi-Fi Alliance test bed. To date Symbol does not have any client devices that are CCX certified. Their mobile WLAN clients use a power save option to conserve battery life. The power save option puts the onboard radio to sleep for a period of time. The sleep time is for most WLAN clients is 100 milliseconds, which is the default time interval of the beacon transmissions, but is a configuration option. There is a second parameter known as a DTIM that figures in the time a WLAN client maybe at sleep. The power save client may use the DTIM value, times the beacon interval, to create a period of sleep time. If the DTIM value is 3, then the sleep time will be 300ms for a beacon interval of 100ms. The access point learns if a WLAN client is

configured for power save from the association process of the client. The access points keep track of power save clients. When a power save client is associated to the access point, then the access point does not send packets to that client while the client is sleeping. The access point buffers those packets with the intent of sending them shortly after the DTIM interval. Voice and video queues are more likely to fill up in 100 to 200ms then best effort queues.

There are many models of Symbol mobile handsets that are WLAN capable. Symbol has been manufacturing wireless mobile clients for 20 years and been manufacturing 802.11 WLAN clients for 10 years. They have a parallel experience to Aironet products. Over ten years the WLAN capabilities of their clients have gone from the original 802.11 data rates of 1Mbps and 2Mbps to the current 802.11g specification of 54Mbps. It is important to know what the capabilities of the different models of Symbols handhelds, from both radio performance and security stand point. The radio performance influences the cell size and the cell throughput.

Cisco Assurewave Program

Mobility services within healthcare organizations worldwide has become pervasive. Clinical workflow and therefore patient care can be impacted as a result of the various wireless medical devices increasingly found in healthcare settings. As a response to the concerns of Medical Device Manufacturers and hospitals worldwide, Cisco has introduced the AssureWave program.

AssureWave is a Cisco program that focuses on satisfying customer wireless quality requirements in critical markets, such as healthcare. This program links and expands on product testing conducted within development engineering, regression testing, and systems test groups within Cisco Systems. AssureWave certification marks the successful completion of extensive Wireless LAN Controller and AP testing validating targeted releases. The test networks and accompanying clients, applications, and features used are gleaned from various sources, including the Cisco Technical Assistance Center (TAC), Cisco Sales, and support teams, as well as directly from healthcare customers. This input is critical to make each test network a true reflection of a functioning healthcare customer environment.

The goal of AssureWave is straight-forward: improve the quality of the release through direct customer involvement and testing. AssureWave provides testing coverage for critical feature areas as required by customers. It complements internal product testing efforts with customer-specific testing to certify functionality. Most importantly, AssureWave delivers on the quality commitment provided to Cisco customers.

Assurewave involves testing select code releases and feature sets, as well as using our Healthcare partners actual equipment in a realistic healthcare environment while using Wireless LAN Controllers (WLC) suited to handle the expected traffic and deployment expectations. Coverage is currently in place for the 210x, 440x, ISR Integrated Wireless Service Module (WLCM), Catalyst 3750 Integrated Wireless Switch, and the Catalyst 6500 Series Wireless LAN Services Module (WiSM). This combination of features, hardware, and image set is tested in a laboratory environment that simulates the healthcare network environment. Cisco updates its testing with best practices guidelines as well as topologies and configurations provided by customers deploying the Cisco Wireless LAN Controller Series in their environment. Test results are unique to technologies covered and actual scenarios in which they were tested.

One of the key elements about the AssureWave program is that we have established partnerships with major device and application vendors who perform extensive testing at their own facilities to ensure broader interoperability with our ongoing releases. These partners, like Cisco, are technology champions and the well known leaders in their respective healthcare focus.

AssureWave test documentation stipulates that the tests either Pass, Pass with Exception, or Fail. Testing schedules are based on code quality, not a date target.

- **Pass**—The underlying assumption for certifying and publishing an AssureWave release is that testing passed, because all individual tests passed. Failure of any test has to be properly resolved, closed, or determined by the AssureWave engineering team to be a non-impacting defect.
- **Pass with Exception**—Indicates that open caveats exist that may cause issues in your network with specific devices or deployment models. These caveats and possible workarounds are documented in the release notes for the firmware used during the evaluation. The AssureWave engineering team opens Cisco DDTS tickets to track issues discovered in Cisco products, as well as documenting possible workarounds.
- **Fail**—If a given test fails, and the impact on Cisco’s customer base is decided to be broad enough, the entire release fails. Failed releases are not certified nor documented. If a test fails, and the impact to the customer base is identified to be minor, the release may still be certified, with DDTSs noted so that customers can review the testing to see if they are impacted.



CHAPTER 4

Symmetric Mobility Tunneling Deployment Guide

Introduction

Cisco Wireless LAN Controllers allow users to roam transparently across all access points in the network. Clients roam seamlessly and maintain IP addressing and session state, even where controllers reside across routed boundaries from one another.

This Layer 3 mobility functionality was designed to deliver traffic with as little added latency as possible, with the simultaneous aim of allowing clients to roam across wireless networks of all scales without requiring any alteration to wired network configuration. This allows controllers to be placed anywhere in the network and as clients roam from AP to AP across these controllers, client connectivity persists, IP addresses remain unchanged, and session state is preserved.

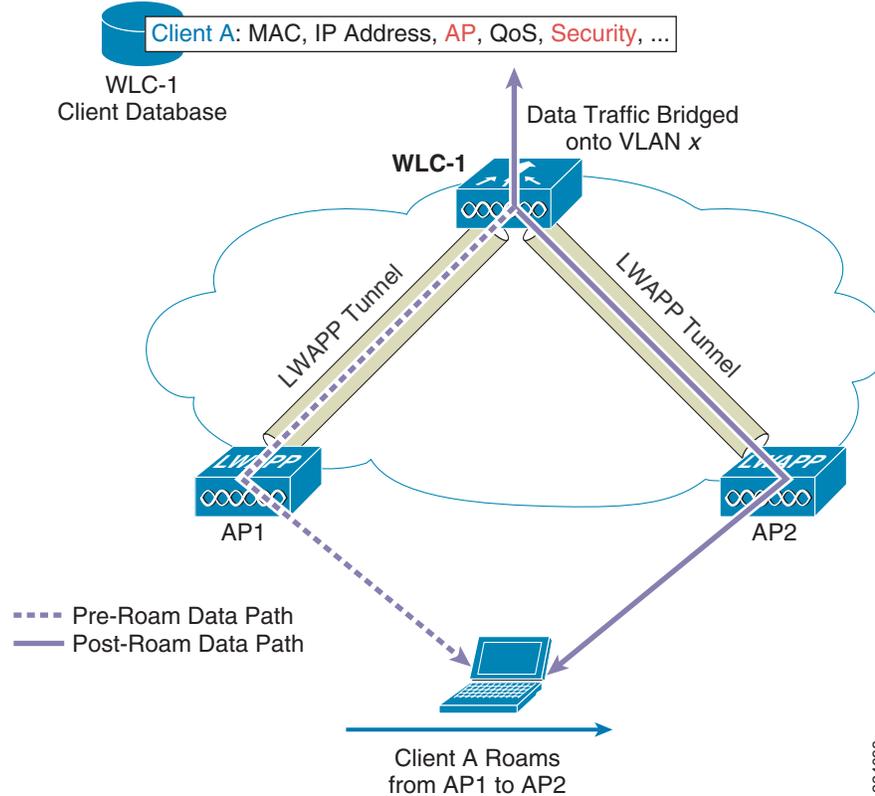
The way this seamless Layer 3 mobility capability was achieved within Cisco's Unified Wireless Network architecture was with an asymmetric traffic pattern whereby a roamed client's egress traffic terminates on the new, foreign subnet (sourced from its original IP address). The client's ingress traffic is routed according to the original IP address that it has maintained, which means that it arrives at its original, anchor controller, is then tunneled to the new, foreign controller, and then delivered to the roamed client.

This Layer 3 mobility operation predominately finds itself working flawlessly in most environments. It is only in networks where traffic is not allowed to be sourced on subnets to which it is not native (due to RPF checks or firewall rules, to name a few curtailing factors), that this asymmetric mobility tunneling between controllers does not function properly. Cisco's new Symmetric Mobility Tunneling feature is aimed squarely at correcting this issue.

Background on Mobility in Cisco's Unified Wireless Network

When a wireless client associates and authenticates to an AP, the AP's joined WLC places an entry for that client in its client database. This entry includes the client's MAC and IP addresses, security context and associations, QoS context, WLAN, and associated AP. The WLC uses this information to forward frames to and from the wireless client. [Figure 4-1](#) depicts what happens when the wireless client roams from one AP to another when both APs are joined to the same WLC.

Figure 4-1 Wireless Client Roaming When Access Points Are Joined to the Same WLC

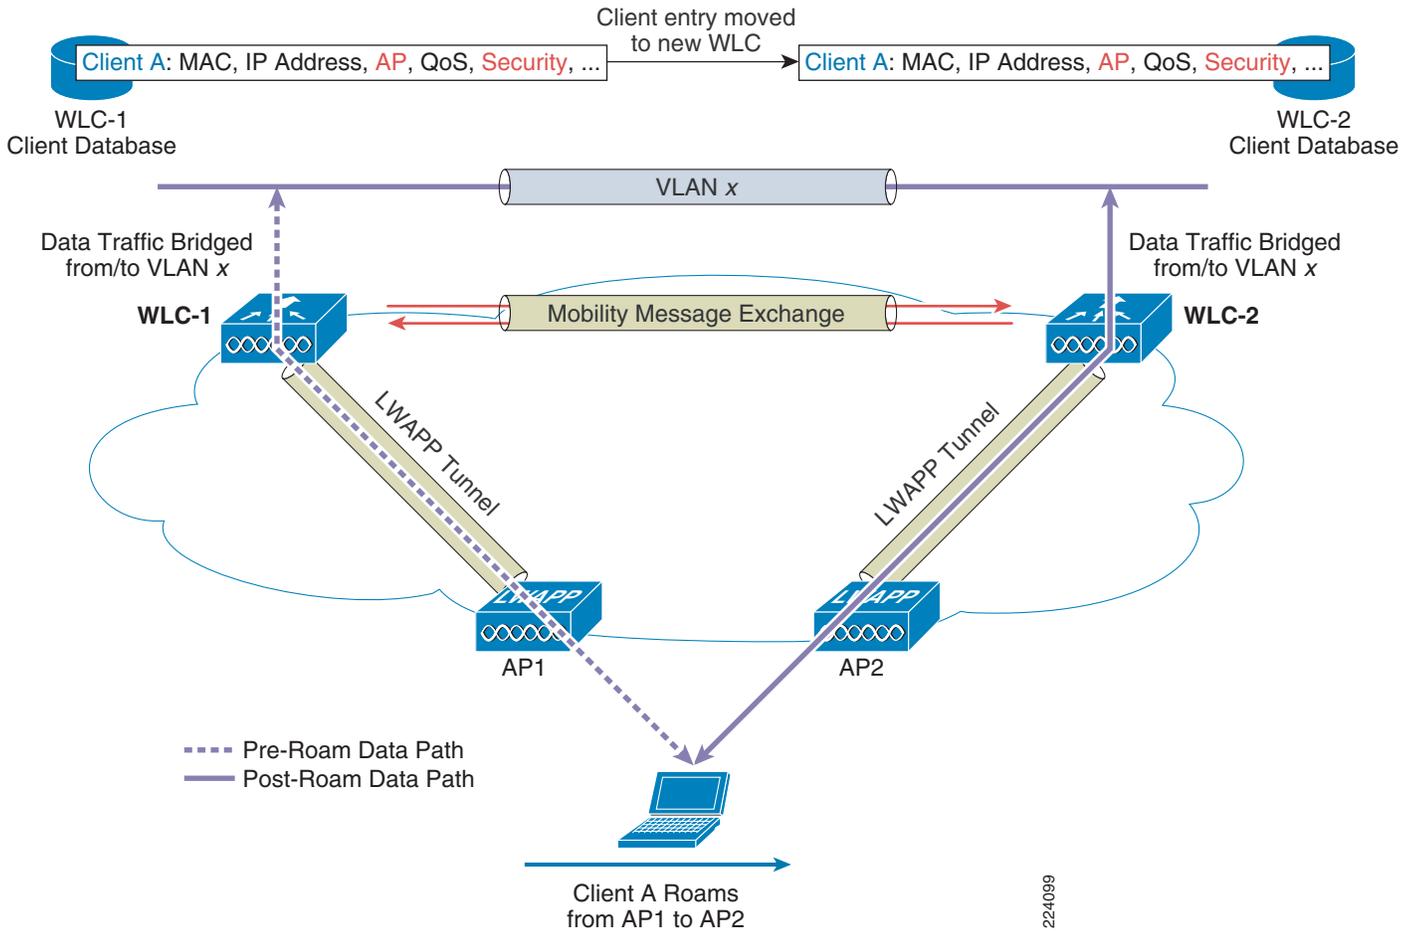


224098

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

Now, consider what happens when a client roams from an AP joined to one WLC and an AP joined to a different WLC. [Figure 4-2](#) illustrates an inter-controller roam in the event of a “Layer 2” roam, whereby the participating controllers are terminating the given WLAN’s traffic on the same subnet.

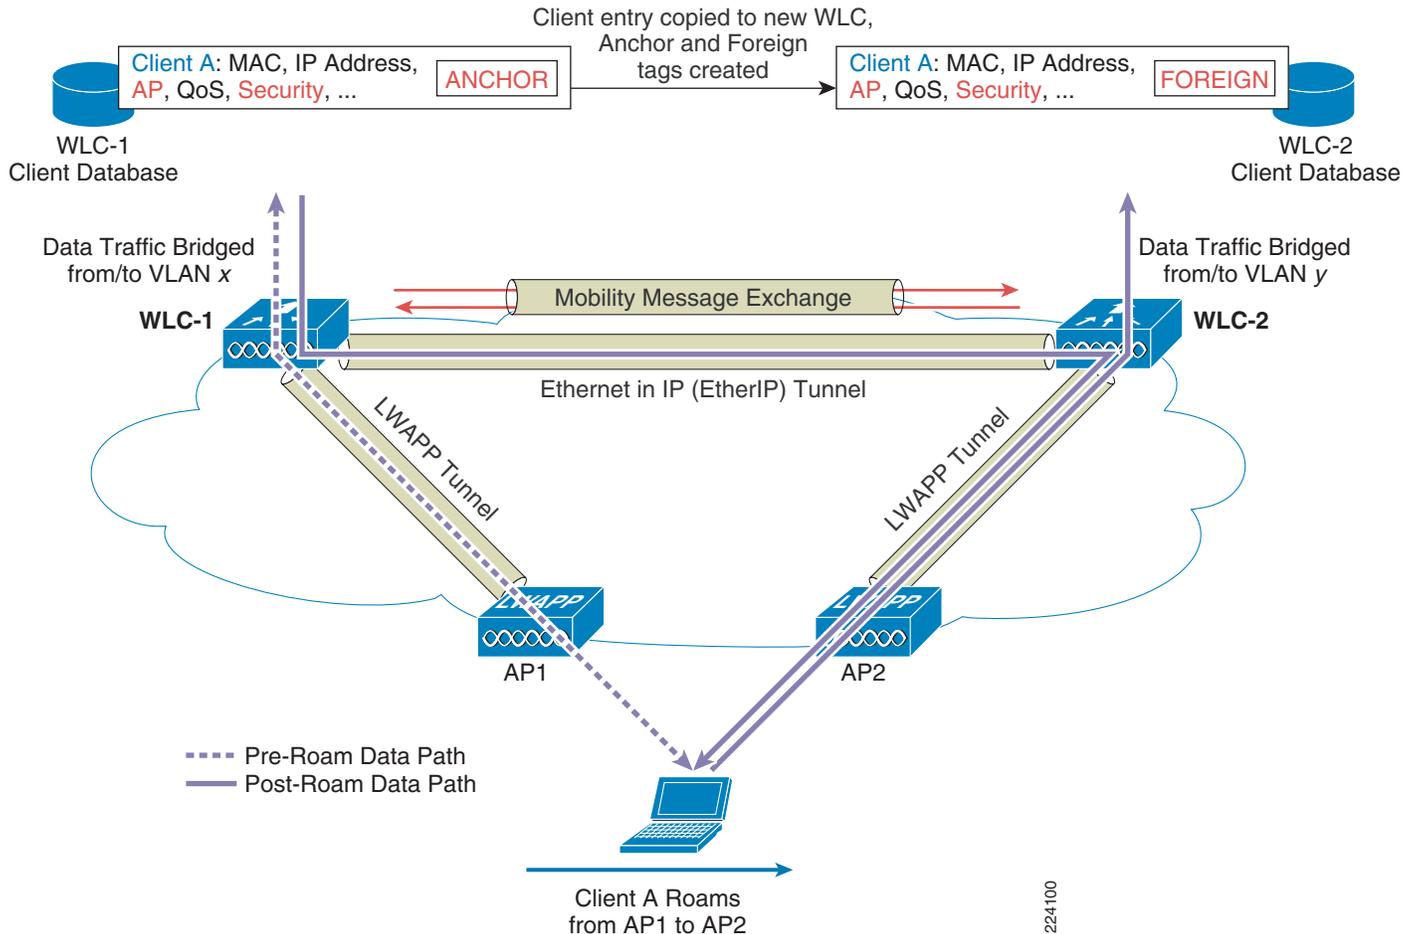
Figure 4-2 Wireless Client Roaming When Access Points Are Joined to Different WLCs



As illustrated, a Layer 2 roam occurs when the controllers bridge the WLAN traffic on and off the same VLAN and the same IP subnet. When the client re-associates to an AP connected to a new WLC, the new WLC exchanges mobility messages (via UDP port 16666, or 16667, if controllers are configured to secure these messages with AES) with the original WLC and the client database entry is moved to the new WLC. New security context and associations are established, if necessary, and the client database entry is updated for the new AP. All of this is transparent to the end user.

Figure 4-3 illustrates an inter-controller roam in the event of a “Layer 3” roam, whereby the participating controllers are not terminating the given WLAN’s traffic on the same subnet.

Figure 4-3 Inter-Controller Roaming With a "Layer 3" Roam



In Figure 4-3, a Layer 3 roam occurs when the controllers bridge the WLAN on and off different VLANs and IP subnets. The inter-controller roaming is similar to Layer 2 roaming in that the WLCs exchange mobility messages upon a client roaming. However, instead of moving the client's entry to the new controller's client database, the original WLAN controller marks the client with an "Anchor" entry in its own client database. The database entry is copied to the new controller client database and marked with a "Foreign" entry in the new WLC. The roam is still transparent to the wireless client and the wireless client maintains its original IP address. Security credentials and context are reestablished if necessary.

After a Layer 3 roam, data to and from the wireless client flows in an asymmetric traffic path. Traffic from the client to the network is forwarded directly into the network by the foreign WLC. Traffic to the client arrives at the Anchor WLC, which forwards the traffic to the Foreign WLC in an Ethernet-in-IP (EtherIP, defined in IETF RFC 3378) tunnel. The Foreign WLC then forwards the data to the client.

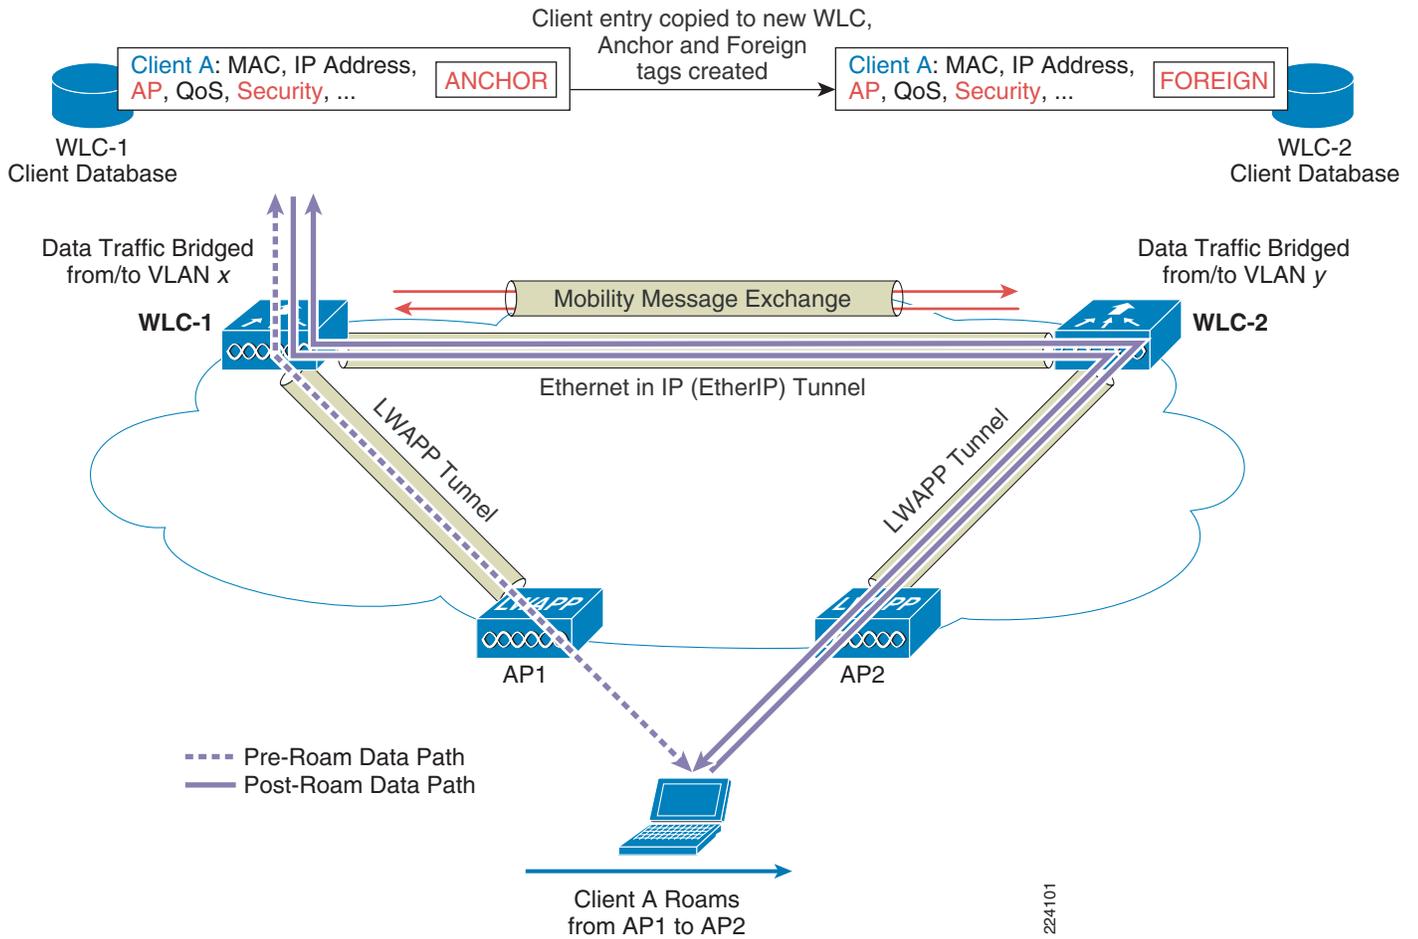
**Note**

If a wireless client roams to a new Foreign WLC, the client database entry is moved from the original Foreign WLC to the new Foreign WLC, but the original Anchor WLC is always maintained.

Symmetric Mobility Tunneling Operation

The new Symmetric Mobility Tunneling feature in the WLC 4.1 release allows both roamed clients' ingress and egress traffic to be tunneled to and from the anchor controller. This means that roamed clients reside logically in their anchor controller and traffic patterns between the anchor and foreign controllers operates fully as a point-to-point symmetric tunnel. The only difference in operation between regular, asymmetric mobility tunneling and this new symmetric traffic flow is that the upstream traffic from roamed clients is not forwarded to the destination by the foreign controller, but is instead first tunneled to the anchor controller, where delivery to the network occurs, as illustrated in Figure 4-4.

Figure 4-4 Symmetric Mobility Tunneling



This feature allows the underlying wired network architecture to remain fully unchanged when such security features as reverse path forwarding/filtering (RPF) checking is enabled on intermediary Layer 3 interfaces or when firewall rules prevent such operation between controllers configured in a mobility group (a cluster of controllers between which roaming is desired).

Mobility Configuration

The first step to configure Symmetric Mobility Tunneling is to verify that all controllers between which seamless roaming must occur are properly configured for mobility operations. Once basic mobility is configured and verified, Symmetric Mobility Tunneling may be enabled.

Mobility configurations can be made through WCS or through the controller's graphical and command line interfaces (though only one configuration interface needs be employed for each given configuration step). The following instructions indicate how to configure basic, asymmetric mobility tunneling first via WCS and then via the controller GUI.

Regular, Asymmetric Mobility Configuration in WCS

From the Configure tab, select Controllers and click the controller of choice. On the left, select the System heading and click the General subheading. Next, make sure all controllers share the same Mobility Domain Name (sometimes referred to as the Mobility Group Name).

Figure 4-5 Controllers Share the Same Mobility Domain Name

The screenshot displays the Cisco WCS configuration interface for a controller. The left sidebar shows the navigation menu with 'System' selected. The main content area shows the 'General' configuration page for the controller at IP 171.71.128.75. The 'Default Mobility Domain Name' is set to 'SJCwireless'. Other settings include 802.3x Flow Control Mode (Disable), 802.3 Bridging (Disable), LWAPP Transport Mode (Layer3), Current LWAPP Operating Mode (Layer3), Ethernet Multicast Support (Multicast), Multicast Group Address (239.255.1.1), Aggressive Load Balancing (Enable), Peer to Peer Blocking Mode (Disable), Over Air Provision AP Mode (Disable), AP Fallback (Enable), Apple Talk Bridging (Disable), Fast SSID change (Disable), Master Controller Mode (Disable), Wireless Management (Enable), Link Aggregation (Disable), and Symmetric Tunneling Mode on next reboot (Disable). The Cisco Discovery Protocol (CDP) section shows CDP on controller (Enable), Global CDP on APs (Enable), Refresh-time Interval (60 seconds), Holdtime (180 seconds), and CDP Advertisement Version (v1). An Alarm Summary table is visible in the bottom left corner.

Category	Count	Severity	Count
Rogue AP	0	0	1079
Coverage	0	0	3
Hole	0	0	0
Security	156	0	47
Controllers	0	1	2
Access Points	162	0	32
Location	0	1	2
Mesh Links	0	0	0
WCS	1	0	0

This can be viewed easily for all controllers at the main controller listing (by going to Configure | Controllers). Once the Mobility Group Name is properly configured, select the Mobility Groups subheading under the System heading. The selected controller’s Mobility List is then shown.

Figure 4-6 Selected Controller’s Mobility List

The screenshot shows the Cisco Wireless Control System (WCS) interface. The breadcrumb path is: 171.71.128.75 > Mobility Group 'SJCwireless' > Group Members. A table lists the group members:

Controller Name	Member MAC Address	Member IP Address	Group Name
<input type="checkbox"/> SJC 14 LWAPP2	00:0b:85:40:4a:60	171.71.128.78	SJCwireless
<input type="checkbox"/> SJC 14 LWAPP1	00:0b:85:46:55:60	171.71.128.75	(Local)

An 'Alarm Summary' table is located at the bottom left of the interface:

Alarm Type	Count	Severity
Rogue AP	0	0
Coverage Hole	0	0
Security	156	0
Controllers	0	1
Access Points	169	0
Location	0	1
Mesh Links	0	0
WCS	1	0

224103

To add members to the list, in the upper right corner, select Add Group Members... from the drop down menu and click Go. All the controllers WCS is managing that are not in the individual controller’s list are displayed. Select the checkbox to the left of desired controllers and click Save.

Figure 4-7 Adding Group Members

The screenshot shows the same WCS interface as Figure 4-6, but with a dropdown menu open in the upper right corner. The dropdown menu contains the following options:

- Select a command --
- Select a command --
- Add Group Members...
- Delete Group Members

224163

**Note**

If this same operation needs to be performed across multiple controllers, it may be prudent to use WCS's controller templates feature to simultaneously push an identical configuration to a group of controllers. This can be done via Configure | Controller Templates.

Figure 4-8 WCS's Controller Templates

The screenshot shows the Cisco WCS interface for the 'SJCwireless' group. The breadcrumb path is '171.71.128.75 > Mobility Group 'SJCwireless' > Group Members'. A table lists the group members:

<input type="checkbox"/>	Controller Name	Member MAC Address	Member IP Address	Group Name
<input checked="" type="checkbox"/>	SJC 14 LWAPP2	00:0b:85:40:4a:60	171.71.128.78	SJCwireless
<input type="checkbox"/>	SJC 14 LWAPP1	00:0b:85:46:55:60	171.71.128.75	(Local)

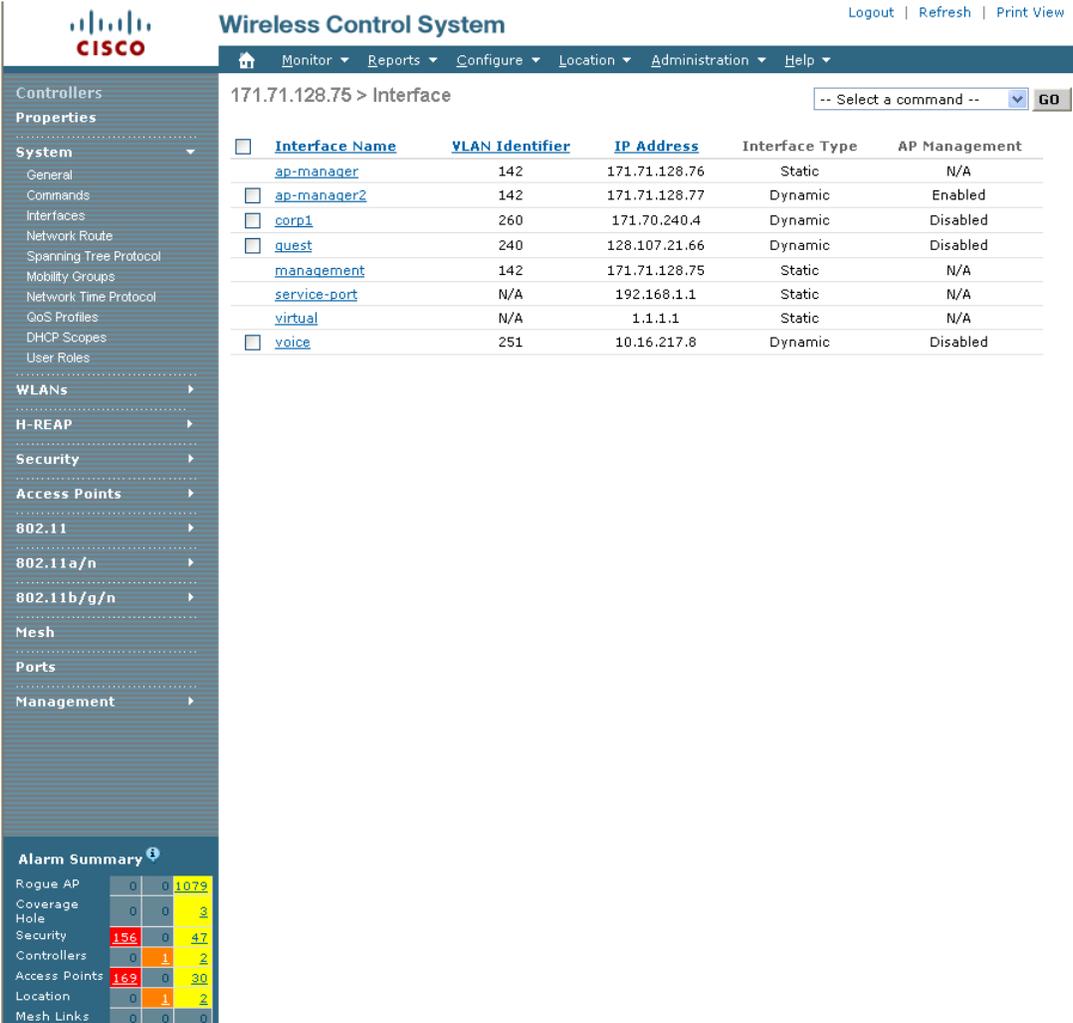
Below the table is an 'Alarm Summary' section with the following data:

Alarm Category	Count	Count	Count
Rogue AP	0	0	1079
Coverage Hole	0	0	3
Security	196	0	47
Controllers	0	1	2
Access Points	169	0	30
Location	0	1	2
Mesh Links	0	0	0

224104

Next, ensure that all controllers share the same Virtual Interface address by selecting Interfaces on the left under the System heading.

Figure 4-9 Controllers Share Virtual Interface Address



The screenshot shows the Cisco Wireless Control System GUI. The breadcrumb path is 171.71.128.75 > Interface. A search bar contains "-- Select a command --" and a "GO" button. A table lists interfaces with columns for Interface Name, VLAN Identifier, IP Address, Interface Type, and AP Management. The 'virtual' interface is highlighted in blue.

Interface Name	VLAN Identifier	IP Address	Interface Type	AP Management
ap-manager	142	171.71.128.76	Static	N/A
<input type="checkbox"/> ap-manager2	142	171.71.128.77	Dynamic	Enabled
<input type="checkbox"/> corp1	260	171.70.240.4	Dynamic	Disabled
<input type="checkbox"/> quest	240	128.107.21.66	Dynamic	Disabled
management	142	171.71.128.75	Static	N/A
service-port	N/A	192.168.1.1	Static	N/A
virtual	N/A	1.1.1.1	Static	N/A
<input type="checkbox"/> voice	251	10.16.217.8	Dynamic	Disabled

At the bottom left, an Alarm Summary table is visible:

Alarm Category	Count	Severity
Rogue AP	0	0
Coverage Hole	0	0
Security	156	0
Controllers	0	1
Access Points	169	0
Location	0	1
Mesh Links	0	0

224105

If it is necessary to change this value to set all controllers to the same address, simply click the virtual link under the Interface Name column heading. Change the address and click Save.

**Note**

Remember that this value must be a non-routed address and must be identical across all controllers in the mobility group.

Return to the main list of controllers by clicking Configure | Controllers and make sure all other necessary controllers are properly configured with identical Mobility Group Names and have all other controllers in the group in their Mobility Lists. Also, ensure that all controllers share the same Virtual Interface address.

Regular, Asymmetric Mobility Configuration in the Controller GUI

Select the Controller tab at the top of the screen. The General heading is shown initially. Ensure that the Default Mobility Domain Name value is consistent across all necessary controllers.

Figure 4-10 Domain Name Value Consistent Across Controllers

The screenshot shows the Cisco Controller configuration interface. The 'CONTROLLER' tab is selected, and the 'General' configuration page is displayed. The 'Domain Name' field is set to 'mobile-1'. Other configuration options include Name (Cisco_ff:e7:eb), Control Mode (Disabled), Transport Mode (Layer 3), LAG Mode on next reboot (Enabled), Multicast Mode (Disabled), Broadcast Forwarding (Disabled), Aggressive Load Balancing (Disabled), Over The Air Provisioning of AP (Disabled), AP Fallback (Enabled), Apple Talk Bridging (Disabled), Fast SSID change (Disabled), Default Mobility (mobile-1), RF-Network Name (mobile-1), User Idle Timeout (300 seconds), ARP Timeout (300 seconds), Web Radius Authentication (PAP), 802.3 Bridging (Disabled), Operating Environment (Commercial), and Internal Temp Alarm Limits (0 to 65 C).

Next, make sure all controllers have each other's MAC and IP addresses input into all other controllers' Mobility Lists. View this list by selecting the Mobility Management heading under the Controller tab and then clicking the Mobility Groups subheading.

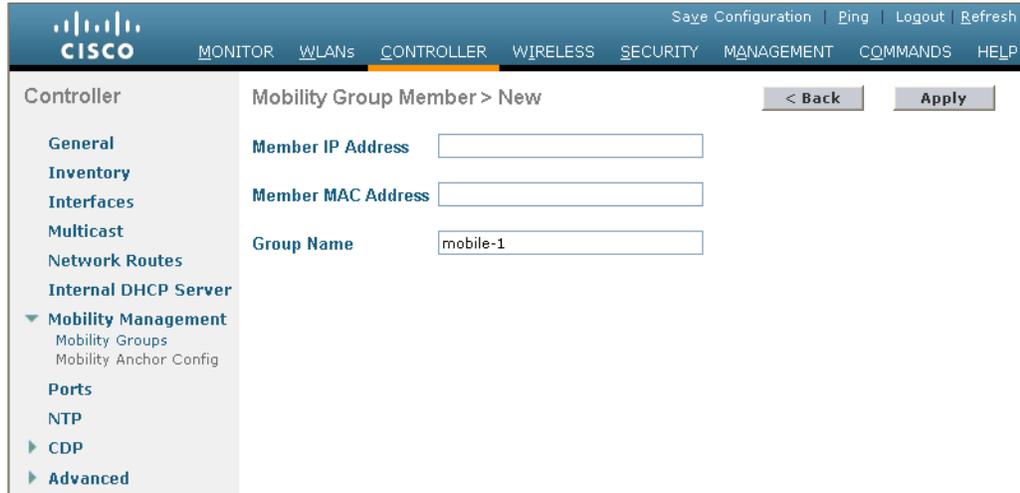
Figure 4-11 Controller Information in All Mobility Lists

The screenshot shows the Cisco Controller configuration interface. The 'CONTROLLER' tab is selected, and the 'Mobility Management' subheading is expanded to show 'Static Mobility Group Members'. The table displays the following information:

MAC Address	IP Address	Group Name
00:11:92:ff:e7:e0	172.20.225.147	(Local)

A single controller may be added by clicking the New... box in the upper right corner, inputting a single controller's information, and then clicking Apply.

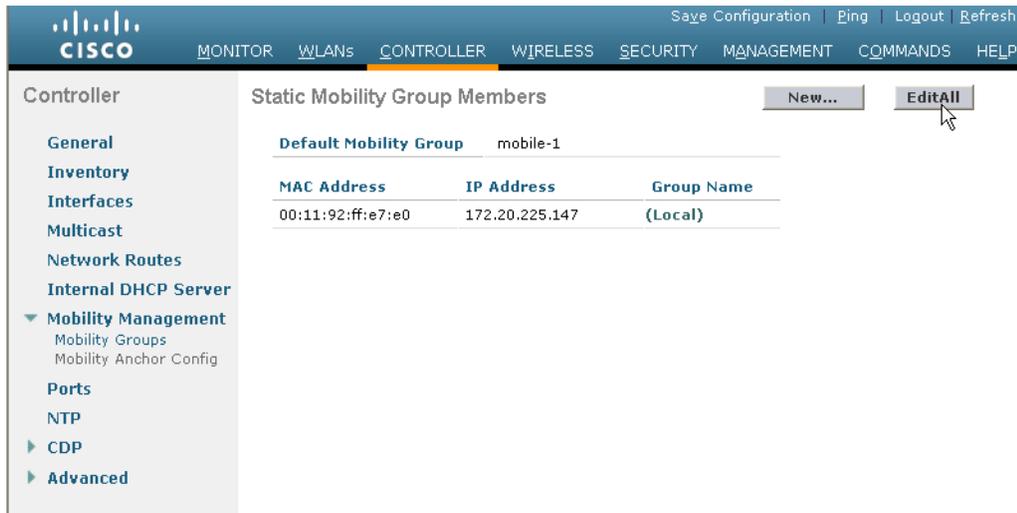
Figure 4-12 Adding a Single Controller



224108

Multiple controllers may be added at once by selecting Edit All. Next, ensure all WLCs have the same Virtual Interface address by selecting the Interfaces heading under the Controller tab.

Figure 4-13 Adding Multiple Controllers



224109

If changes are necessary, click virtual under the column labeled Interface Name and make the necessary changes. Select Apply.

All controllers are now properly configured for regular mobility.

Verifying Regular, Asymmetric Mobility Operation

Verify that regular mobility is operational.

**Note**

To properly trigger the asymmetric mobility tunneling feature (as well as the new Symmetric Mobility Tunneling feature), controllers must be across routed boundaries. If controllers are on the same subnet, then mobility events do not invoke the plumbing of this tunnel because the client record is simply moved to the next controller and traffic flows natively to and from that new controller (see [Background on Mobility in Cisco's Unified Wireless Network](#) for a more in-depth discussion of mobility operations).

The previous configuration steps should be followed to ensure correct configuration. Mobility configuration can be seen easily via the controller CLI.

```
(Cisco Controller) >show mobility summary
```

```
Symmetric Mobility Tunneling (current) ..... Disabled
Symmetric Mobility Tunneling (after reboot) ..... Enabled
Mobility Protocol Port..... 16666
Mobility Security Mode..... Disabled
Default Mobility Domain..... test
Mobility Keepalive interval..... 10
Mobility Keepalive count..... 3
Mobility Group members configured..... 2
```

```
Controllers configured in the Mobility Group
```

MAC Address	IP Address	Group Name	Status
00:16:9d:ca:dc:c0	10.10.10.10	<local>	Up
00:19:07:24:12:e0	20.20.20.20	test	Up

The simplest indicators that the network is properly configured for mobility are two ping variants that can be run between controllers. To verify that configuration is sound and the intermediary network is properly forwarding the necessary traffic, both the `eping` and `mping` commands may be run through the controller command line. The following command can be used to test the operation of the EtherIP data tunnel between controllers:

```
(Cisco Controller) >eping [peer controller's management interface IP address]
```

Similarly, the operation of the UDP port 16666/16667 inter-controller management path can be tested by the following command:

```
(Cisco Controller) >mping [peer controller's management interface IP address]
```

Once mobility has been verified to be properly configured and operational, the wireless network can then be configured for Symmetric Mobility Tunneling.

**Note**

All controllers in the mobility group **must** be configured for Symmetric Mobility Tunneling for the feature to work properly.

Symmetric Mobility Configuration in WCS

Select Configure | Controllers and then select the controller of choice. On the left hand side, select the System heading and then click the General subheading. Next to 'Symmetric Mobility Tunneling Mode on next reboot, choose Enable from the dropdown menu. At the bottom of the page, click Save.

Figure 4-14 Symmetric Mobility Configuration in WCS

The screenshot shows the Cisco WCS interface for controller 171.71.128.75. The left sidebar contains a navigation tree with categories like Properties, System, WLANs, H-REAP, Security, Access Points, Mesh, Ports, and Management. The main area displays the 'General' configuration page for the controller. A dropdown menu is open for the 'Symmetric Tunneling Mode on next reboot' setting, showing 'Enable' and 'Disable' options. A tooltip indicates '(Mode is currently Disabled)'. Below this, the 'Cisco Discovery Protocol' section is visible with various settings like 'CDP on controller' and 'Global CDP on APs' set to 'Enable'. An 'Alarm Summary' table is located at the bottom left of the configuration area.

Category	Count	Count	Count
Rogue AP	0	0	107
Coverage Hole	0	0	3
Security	156	0	47
Controllers	0	1	2
Access Points	169	0	33
Location	0	1	2
Mesh Links	0	0	0

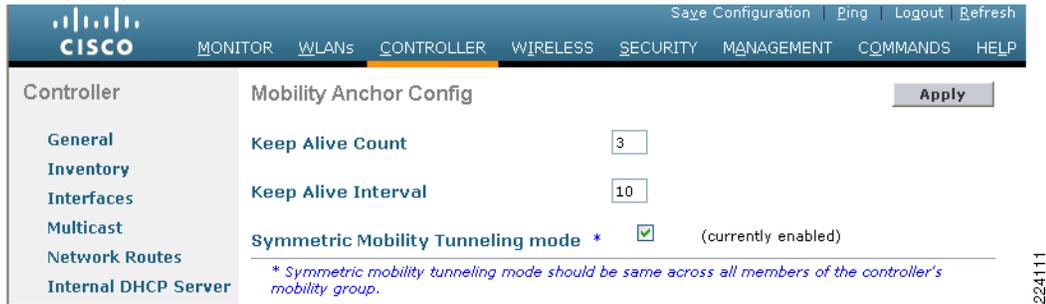
Note

If this same operation needs to be performed across multiple controllers, it may be prudent to use WCS’s controller templates feature to simultaneously push an identical configuration to a group of controllers. This can be done via Configure | Controller Templates.

Symmetric Mobility Configuration in the Controller GUI

To configure Symmetric Mobility Tunneling in the WLC GUI, go to Controller | Mobility Management and then select the Mobility Anchor Config subheading. Check the box next to Symmetric Mobility Tunneling mode and click Apply.

Figure 4-15 Symmetric Mobility Configuration in the Controller GUI



To configure this through the controller's CLI, simply enter the following command:

```
(Cisco Controller) >config mobility symmetric-tunneling enable
```

Upon configuring the controller for Symmetric Mobility Tunneling, the configuration must be saved and the controller rebooted. Each controller in the mobility group must have this operation performed. This can be eased for all controllers through WCS.



Note

Make sure all configurations are saved and the controllers rebooted. Without this step, Symmetric Mobility Tunneling will not work.

Configuration of Symmetric Mobility Tunneling is complete.



APPENDIX **A**

Antenna Recommendations

Antenna Types and Placement Recommendations

There are many different antennas available from Cisco Systems, Inc. The online location of the Cisco Aironet Antennas and Accessories Reference Guide is shown in [Figure A-1](#) in the Wireless category and Antennas sub-category (<http://www.cisco.com/web/psa/products/index.html>).

Figure A-1 Cisco Aironet Antennas and Accessories Reference Guide

CISCO

Solutions | Products & Services | Ordering | Support | Training & Events | Partner Central

HOME
SUPPORT
Product Support

Product Support
Select Your Product

Products | Technology

Search for a specific Product

e.g. Catalyst 6500 Series Switches or 7600 Series Routers

Or Select from a category

- 1 Select a category**
 - Application Networking Services
 - Cisco IOS Software
 - Cisco Interfaces and Modules
 - Interoperability Systems
 - Network Management
 - Optical Networking
 - Routers
 - Security
 - Server Networking and Virtualization
 - Service Exchange
 - Storage Networking
 - Switches
 - TelePresence
 - Universal Gateways and Access Servers
 - Video, Cable and Content Delivery
 - Voice and Unified Communications
 - Wireless**
 -
 - [End-of-Sale and End-of-Life Products](#)
- 2 Select a sub-category**
 - Access Points
 - Antennas**
 - Cisco Mobility Express
 - Client Adapters and Client Software
 - Mobile Wireless
 - Outdoor Wireless
 - Service Control for Mobile Networks
 - Wireless IP Telephony
 - Wireless Integrated Switches and Routers
 - Wireless LAN Controllers**
 - Standalone Controllers
 - Wireless Integrated Routers
 - Wireless Integrated Switches
 - Wireless LAN Management**
 - Autonomous Wireless LAN Management
 - Unified Wireless LAN Management
 - Wireless Security Servers
 - Wireless Spectrum Analyzer
- 3 Select a product**
 - [Cisco Aironet Antennas and Accessories](#)

224049

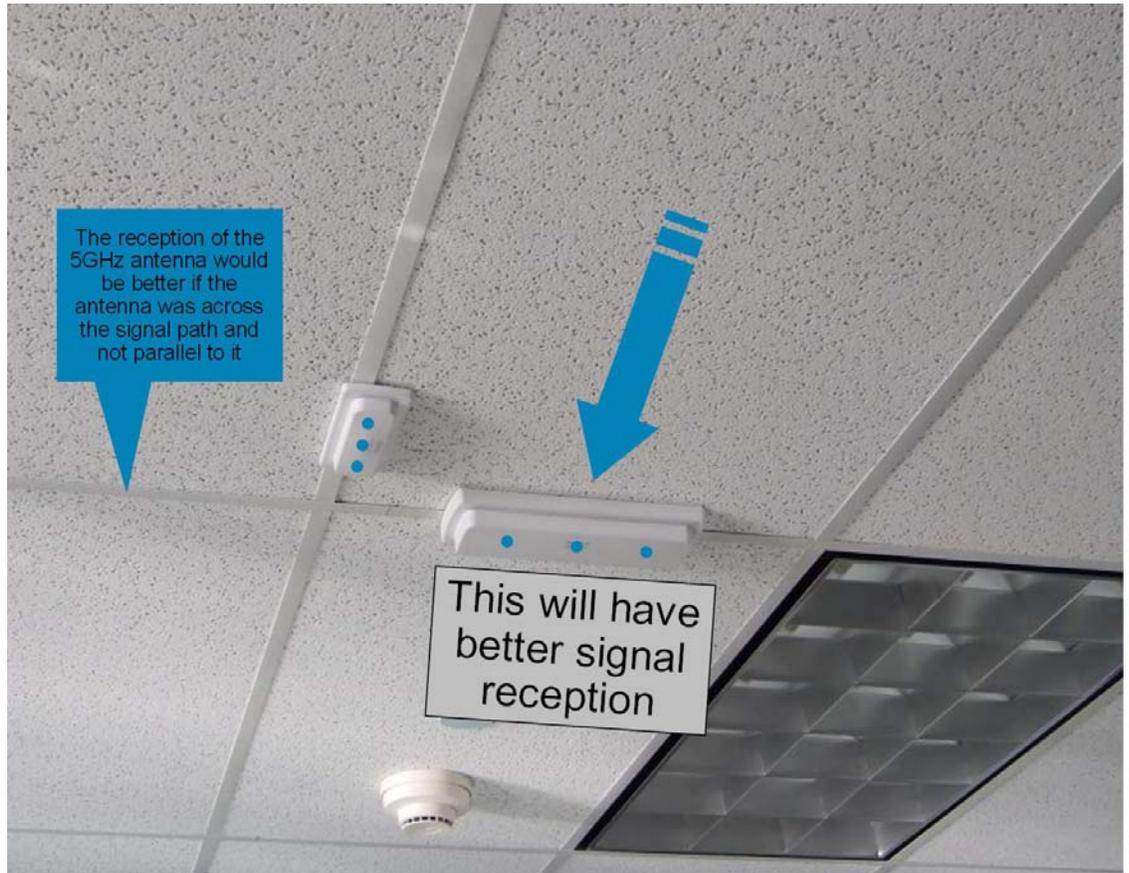
In hospitals it is always recommended that diversity antennas be used for 2.4GHz and 5GHz. If the access point is 802.11n capability, then a MIMO antenna is recommended. The diversity antennas have two receiving elements and the MIMO antennas have three receiving elements. Using an antenna with diversity (two elements) provides twice the likelihood of receiving a reflected/multipath signal over one antenna while MIMO antennas provide three times the likelihood. Both diversity and MIMO antennas reduce retries and improve cell throughput. In some sites the combination of diversity antenna clients and diversity antenna access points has reduced retry counts by 50%.

The AIR-ANT5959 2.4GHz and AIR-ANT5145-R 5GHz are both diversity Omni that mount on the T-Bar just below the ceiling tile. These antennas can be used on the 1240 series access point. The AIR-ANT2430V-R for 2.4GHz and the AIR-ANT5140V-R for 5 GHz are MIMO Omnis that match up to the 1250 series access point. They work well in most areas of the hospital, but there are many other antennas with more specific performance parameters that may be more appropriate in certain hospital locations. There may be installation and antenna placement techniques other than those provided that need to be used. However these guidelines provide insight into the commonly deployed techniques.

**Note**

There may be an occasion when an AP1250 replaces an AP1240. The AP1250 could operate with the antennas used for the AP1240 for a temporary fix. Leave the C-Rx port of the AP1250 unused. Connect the AP1240 antennas to the A-Tx/Rx and B-Tx/Rx ports of the AP1250.

MIMO antennas should be installed with the wider edge of the cover of the access point directed down a hallway. The wider edge should be perpendicular to the hallway. Each of the three receive/transmit elements should have individual exposure to the length of the hallway. It is incorrect to have the third and second elements of the MIMO send signals over the first element down the length of the hallway. This should also be the installation method used in rooms. The wider side of the antenna should be perpendicular to the longer dimension of a room. This is also the recommended placement for diversity Omnis. In [Figure A-2](#) the 2.4 GHz MIMO ANT2430V-R is positioned on the T-Bar so that it has the correct orientation to the length of the hallway. The 5GHz MIMO ANT5140V-R is not recommended and should be moved to the other T-Bar so its orientation matches the 2.4GHz antenna.

Figure A-2 Antenna Placement

The large arrow indicates the direction of the hallway. The 3 small dots on the larger 2.4GHz MIMO antenna and the smaller 5GHz MIMO antenna reference the radiating elements in the antennas.

See [Figure A-3](#) for more information on the recommended AP1250 MIMO antennas.

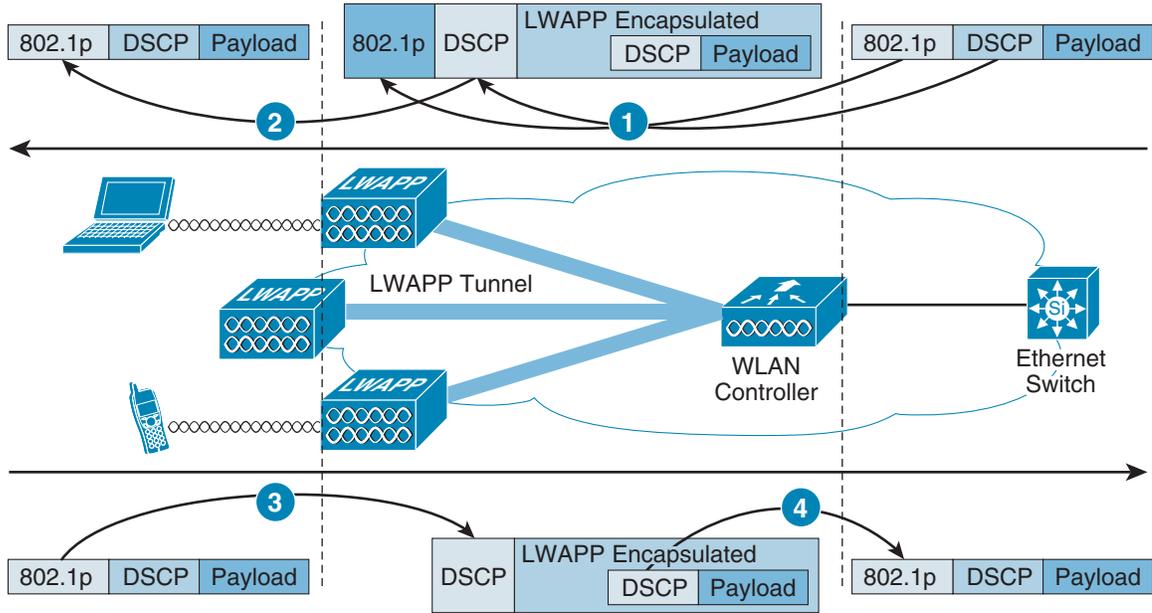
Figure A-3 Recommended AP1250 MIMO Antennas

AP1250	
AIR_ANT2430V-R	AIR_ANT5140V-R
3.0 dBi 2.4GHz Omnidirectional 3 Element diversity	4.0 dBi 5GHz Omnidirectional 3 Element diversity
	

The utilization of channel is influenced by data rates configured on the access points and used by the client devices. The 802.11 data rates of 1Mbps and 2Mbps go back to the original 1997 specification. There are still client devices that require and/or only work with these data rates. For the most part those are data collections devices, such as barcode scanners. If there is no need to support these data rates, then it is highly recommended that these rates be disabled. [Figure A-4](#) shows a hospital that disabled these rates and lowered the channel utilization from over 35% to 5%. This one change fixed several applications. Similar considerations should be made of the lower two data rates for 5GHz 802.11a and 802.11n WLANs.

Figure A-5 Marking of a VoWLAN Transiting from the VoWLAN Client to the WLAPP Controller and Back to the Phone

Markings	802.11e	DSCP	802.1p	DSCP (LWAPP)	DSCP
SCCP	4	CS3	3	CS3	CS3
RTP	6	EF	5	EF	EF



Markings	802.11e	DSCP	DSCP	DSCP (LWAPP)	802.1p
SCCP	4	CS3	CS3	CS3	3
RTP	6	EF	EF	EF	5

224164



APPENDIX **B**

References

- Wi-Fi Alliance
www.wi-fi.org
- Wireless LAN Compliance Status
http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps5861/product_data_sheet0900aecd80537b6a_ps4570_Products_Data_Sheet.html
- Cisco Wireless Control System Configuration Guide 4.2
www.cisco.com/en/US/products/ps6305/tsd_products_support_series_home.html
- Cisco Wireless LAN Controller Configuration Guide 4.2
www.cisco.com/en/US/products/ps6366/products_installation_and_configuration_guides_list.html
- Cisco 2700 Series Wireless Location Appliance Deployment Guide
www.cisco.com/en/US/products/ps6386/tsd_products_support_series_home.html
- Wi-Fi Location Based Services Design Guide 4.1
Voice over Wireless LAN 4.1 Design Guide
Enterprise Mobility 4.1 Design Guide
www.cisco.com/go/srmd
- Cisco Unified Wireless IP Phone 7921G Installation Guides
www.cisco.com/en/US/products/hw/phones/ps379/prod_installation_guides_list.html
- Cisco Unified Wireless IP Phone 7921g Deployment Guides
www.cisco.com/en/US/products/hw/phones/ps379/prod_installation_guides_list.html

Links to Web sites and materials mentioned in this document:

- Wi-Fi Health Download File
http://www.wi-fi.org/files/kc_13_WLAN%20and%20Health_7-20-05.pdf
- The World Health Organization
<http://www.who.int/mediacentre/factsheets/fs304/en/index.html>
- Health Physics Society
<http://hps.org/hpspublications/articles/wirelessnetworks.html>
- Bioelectromagnetics Society
www.bioelectromagnetics.org
- European Bioelectromagnetics Association
www.ebea.org
- Electromagnetic Energy Association
www.elecenergy.com

- Federal Communications Commission
www.fcc.gov/oet/rfsafety
- U.S. Food and Drug Administration
www.fda.gov/cdrh/phones/index.html
<http://www.fda.gov/cdrh/osel/guidance/1618.html> Guidance for WiFi in Medical devices
- ICNIRP (Europe)
www.icnirp.de
- IEEE
www.ieee.org
- IEEE Committee on Man & Radiation
www.seas.upenn.edu/~kfoster/comar.htm
- National Council on Radiation Protection & Measurements
www.ncrp.com
- Health Protection Agency (United Kingdom)
<http://www.hpa.org.uk/radiation/>
- US OSHA
www.osha-slc.gov/SLTC/radiofrequencyradiation/index.html
- Wireless Industry (CTIA)
www.wow-com.com