



Enterprise Mobility 7.3 Design Guide

Cisco Validated Design

Revised: September 27, 2013

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

Customer Order Number:
Text Part Number: OL-14435-01

Cisco Validated Design

The Cisco Validated Design Program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit www.cisco.com/go/validateddesigns.

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

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

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

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

Enterprise Mobility 7.3 Design Guide

© 2013 Cisco Systems, Inc. All rights reserved.



Preface i

- Document Purpose i-i
- Intended Audience i-i
- Document Organization i-i

CHAPTER 1

Cisco Unified Wireless Network Solution Overview 1-1

- WLAN Introduction 1-1
- WLAN Solution Benefits 1-1
- Requirements of WLAN Systems 1-2
- Cisco Unified Wireless Network 1-5

CHAPTER 2

Cisco Unified Wireless Technology and Architecture 2-1

- CAPWAP Overview 2-1
 - Split MAC 2-3
 - Layer 3 Tunnels 2-5
 - WLC Discovery and Selection 2-7
 - CAPWAP AP Reset 2-7
- Core Components 2-8
 - Cisco Wireless LAN Controllers 2-8
 - Cisco Access Points 2-10
 - CAPWAP APs 2-10
 - Cisco Prime Infrastructure 2-11
- Mobility Groups, AP Groups, and RF Groups 2-12
 - Mobility Groups 2-12
 - Mobility Group Definition 2-13
 - Mobility Group Application 2-13
 - Mobility Group—Exceptions 2-13
 - AP Groups 2-14
 - RF Groups 2-15
- Roaming 2-16
- Broadcast and Multicast on the WLC 2-19
 - WLC Broadcast and Multicast Details 2-20
 - DHCP 2-21
 - VideoStream 2-21

Other Broadcast and Multicast Traffic	2-21
Design Considerations	2-22
WLC Location	2-22
Distributed WLC Deployment	2-22
Centralized WLC Deployment	2-23
Centralizing WLCs	2-24
Distributed WLC Network Connectivity	2-25
Traffic Load and Wired Network Performance	2-25
Volume of CAPWAP Control Traffic	2-26
Overhead Introduced by Tunneling	2-26
Traffic Engineering	2-26
AP Connectivity	2-27
Operation and Maintenance	2-27
WLC Discovery	2-27
AP Distribution	2-27

CHAPTER 3

WLAN RF Design Considerations 3-1

RF Basics	3-1
Regulatory Domains	3-2
Operating Frequencies	3-2
2.4 GHz - 802.11b/g/n	3-2
5 GHz - 802.11a/n/ac	3-2
Deployment Considerations	3-3
Understanding the IEEE 802.11 Standards	3-4
Direct Sequence Spread Spectrum (DSSS)	3-5
IEEE 802.11b Direct Sequence (DS) Channels	3-5
IEEE 802.11g	3-6
IEEE 802.11a OFDM Physical Layer	3-7
IEEE 802.11a Channels	3-7
RF Power Terminology	3-7
dB	3-7
dBi	3-8
dBm	3-8
Effective Isotropic Radiated Power (EIRP)	3-8
Planning for RF Deployment	3-9
Different Deployment Types of Overlapping WLAN Coverage	3-10
Data-Only Deployments	3-10
Voice Deployments	3-10
Location-Based Services Deployments	3-12

WLAN Data Rate Requirements	3-13
Data Rates Compared to Coverage Area	3-13
AP Density for Different Data Rates	3-14
Client Density and Throughput Requirements	3-16
WLAN Coverage Requirements	3-17
Power Level and Antenna Choice	3-18
Omni-Directional Antennas	3-18
Patch Antennas	3-19
Dipole Antennas	3-20
Security Policy Requirements	3-21
RF Environment	3-21
RF Deployment Best Practices	3-22
Manually Fine-Tuning WLAN Coverage	3-22
Channel and Data Rate Selection	3-23
Recommendations for Channel Selection	3-23
Manual Channel Selection	3-24
Data Rate Selection	3-25
Data Rate Modes	3-26
Lowest and Highest Mandatory Rate Settings	3-27
Radio Resource Management	3-27
Overview of RRM Operation	3-28
RRM Configuration Settings	3-29
Sample 'show ap auto-rf' Command Output	3-32
Dynamic Channel Assignment	3-33
Interference Detection and Avoidance	3-33
Dynamic Transmit Power Control	3-34
Coverage Hole Detection and Correction	3-34
Client and Network Load Balancing	3-34

CHAPTER 4

Cisco Unified Wireless Network Architecture —Base Security Features 4-1

Secure Wireless Topology	4-1
WLAN Security Mechanisms	4-2
Cisco Wired Equivalent Privacy (WEP) Extensions	4-2
Wi-Fi Protected Access (WPA)	4-2
Wi-Fi Protected Access 2 (WPA2)	4-3
802.1X	4-3
Authentication and Encryption	4-3
Extensible Authentication Protocol	4-4
Authentication	4-5

Suplicants	4-5
Authenticator	4-6
Authentication Server	4-7
Encryption	4-8
TKIP Encryption	4-8
AES Encryption	4-9
Four-Way Handshake	4-10
Proactive Key Caching and CCKM	4-11
Cisco Unified Wireless Network Architecture	4-13
CAPWAP Features	4-14
Important Points to Remember	4-14
Cisco Unified Wireless Network Security Features	4-15
Enhanced WLAN Security Options	4-15
Local EAP Authentication	4-17
ACL and Firewall Features	4-18
DHCP and ARP Protection	4-19
Peer-to-Peer Blocking	4-19
Wireless IDS	4-20
Cisco Adaptive Wireless Intrusion Prevention System	4-21
Dedicated Monitor Mode versus ELM	4-22
On-Channel and Off-Channel Performance	4-22
ELM Across WAN Links	4-23
CleanAir Integration	4-23
ELM wIPS Alarm Flow	4-23
Client Exclusion	4-24
Rogue AP	4-25
Air/RF Detection	4-25
Location	4-26
Wire Detection	4-26
Switch Port Tracing	4-27
Rogue AP Containment	4-27
Management Frame Protection	4-27
Client Management Frame Protection	4-29
Management System Security Features	4-30
Configuration Verification	4-30
Alarms and Reports	4-31
Architecture Integration	4-31
Cisco Integrated Security Features	4-32
Types of Attacks	4-32

MAC Flooding Attack	4-32
DHCP Rogue Server Attack	4-33
DHCP Starvation Attack	4-33
ARP Spoofing-based Man-In-the-Middle Attack	4-33
IP Spoofing Attack	4-33
CISF for Wireless Deployment Topologies	4-34
Using Port Security to Mitigate a MAC Flooding Attack	4-35
Port Security in a Wireless Network	4-35
Effectiveness of Port Security	4-36
Using Port Security to Mitigate a DHCP Starvation Attack	4-36
Wireless DHCP Starvation Attack	4-36
DHCP Snooping to Mitigate a Rogue DHCP Server Attack	4-37
DHCP Snooping for Wireless Access	4-37
Effectiveness of DHCP Snooping	4-38
Dynamic ARP Inspection to Mitigate a Man-in-the-Middle Attack	4-38
DAI for Wireless Access	4-38
Effectiveness of DAI	4-39
Using IP Source Guard to Mitigate IP and MAC Spoofing	4-40
IP Source Guard for Wireless Access	4-40
Effectiveness of IP Source Guard	4-41
Summary of Target Attacks	4-42

CHAPTER 5

Cisco Unified Wireless QoS 5-1

QoS Overview	5-1
Wireless QoS Deployment Schemes	5-1
QoS Parameters	5-2
Radio Upstream and Downstream QoS	5-3
QoS and Network Performance	5-4
802.11 Distributed Coordination Function	5-4
Interframe Spaces	5-5
Random Backoff	5-5
aCWmin, aCWmax, and Retries	5-6
Wi-Fi Multimedia	5-7
WMM Access	5-7
WMM Classification	5-8
WMM Queues	5-9
Enhanced Distributed Channel Access	5-10
Unscheduled-Automatic Power-save Delivery	5-12
TSPEC Admission Control	5-14

Advanced QoS Features for WLAN Infrastructure	5-16
QoS Profiles	5-16
WMM Policy	5-18
Voice over IP Phones	5-19
Admission Control Parameters	5-20
Impact of TSpec Admission Control	5-22
802.11e, 802.1P and DSCP Mapping	5-23
QoS Baseline Priority Mapping	5-24
Deploying QoS Features on CAPWAP-based APs	5-25
WAN QoS and FlexConnect	5-25
Guidelines for Deploying Wireless QoS	5-26
QoS LAN Switch Configuration Example	5-26
AP Switch Configuration	5-26
WLC Switch Configuration	5-26
Traffic Shaping, Over the Air QoS, and WMM Clients	5-27
WLAN Voice and Cisco Phones	5-27
CAPWAP over WAN Connections	5-27
CAPWAP Traffic Classification	5-28
CAPWAP Control Traffic	5-28
CAPWAP 802.11 Traffic	5-29
Classification Considerations	5-29
Router Configuration Examples	5-30

CHAPTER 6
Cisco Unified Wireless Multicast Design 6-1

Introduction	6-1
Overview of Multicast Forwarding	6-1
Wireless Multicast Roaming	6-3
Asymmetric Multicast Tunneling	6-3
Multicast Enabled Networks	6-4
CAPWAP Multicast Reserved Ports and Addresses	6-4
Enabling Multicast Forwarding on the Controller	6-5
CLI Commands to Enable Ethernet Multicast Mode	6-5
Multicast Deployment Considerations	6-6
Recommendations for Choosing a CAPWAP Multicast Address	6-6
Fragmentation and CAPWAP Multicast Packets	6-6
All Controllers have the Same CAPWAP Multicast Group	6-7
Controlling Multicast on the WLAN Using Standard Multicast Techniques	6-7
How Controller Placement Impacts Multicast Traffic and Roaming	6-9
Additional Considerations	6-10

CHAPTER 7**FlexConnect 7-1**

Supported Platforms	7-2
FlexConnect Terminology	7-2
Switching Modes	7-2
Local Switched	7-2
Central Switched	7-2
Operation Modes	7-3
FlexConnect States	7-3
Authentication-Central/Switch-Central	7-3
Authentication Down/Switching Down	7-3
Authentication-Central/Switch-Local	7-4
Authentication-Down/Switch-Local	7-4
Authentication-local/switch-local	7-5
Applications	7-5
Branch Wireless Connectivity	7-5
Branch Guest Access	7-6
Public WLAN Hotspot	7-6
Deployment Considerations	7-7
WAN Link	7-7
Roaming	7-8
Radio Resource Management	7-8
Location Services	7-9
QoS Considerations	7-9
General Deployment Considerations	7-9
FlexConnect Solution	7-10
Advantages of Centralizing Access Point Control Traffic	7-10
Advantages of Distributing Client Data Traffic	7-10
Central Client Data Traffic	7-10
Cisco Flex 7500 Series Cloud Controller	7-11
Modes of Operation	7-12
Primary Design Requirements	7-12
Branch Networking Features and Best Practices	7-13
FlexConnect Groups	7-14
Configuring FlexConnect Groups	7-14
CLI Verification	7-16
Local Authentication	7-16
Local EAP	7-17
CCKM/OKC Fast Roaming	7-17
FlexConnect VLAN Override	7-18

FlexConnect VLAN Override Summary	7-18
FlexConnect VLAN Based Central Switching	7-18
FlexConnect VLAN Central Switching Summary	7-19
FlexConnect ACL	7-19
FlexConnect ACL Summary	7-19
FlexConnect ACL Limitations	7-19
FlexConnect Split Tunneling	7-20
Split Tunnel Summary	7-20
Split Tunnel Limitations	7-20
Fault Tolerance	7-21
Fault Tolerance Summary	7-21
Fault Tolerance Limitations	7-21
Peer-to-Peer Blocking	7-21
P2P Summary	7-22
P2P Limitations	7-22
FlexConnect WGB/uWGB Support for Local Switching WLANs	7-22
FlexConnect WGB/uWGB Summary	7-22
FlexConnect WGB/uWGB Limitations	7-23
Guidelines and Limitations	7-23

CHAPTER 8

Cisco Wireless Mesh Networking	8-1
Access Point Roles	8-2
Network Access	8-3
Network Segmentation	8-3
Cisco Indoor Mesh Access Points	8-3
Cisco Outdoor Mesh Access Points	8-4
Cisco Aironet 1552 Mesh Access Point	8-5
Cisco 1522 Mesh Access Point	8-6
Cisco 1524SB Mesh Access Point	8-7
Ethernet Ports	8-7
1550 Series Multiple Power Options	8-8
Cisco Wireless LAN Controllers	8-8
Cisco Prime Infrastructure	8-8
Architecture	8-9
Control and Provisioning of Wireless Access Points	8-9
CAPWAP Discovery on a Mesh Network	8-9
Adaptive Wireless Path Protocol	8-9
Traffic Flow	8-10

Mesh Neighbors, Parents, and Children	8-11
Criteria to Choose the Best Parent	8-12
Ease Calculation	8-12
Parent Decision	8-12
Mesh Deployment Modes	8-13
Wireless Backhaul	8-13
Universal Access	8-13
Point-to-Multipoint Wireless Bridging	8-13
Wireless Backhaul Data Rate	8-14
ClientLink Technology	8-15
Controller Planning	8-16
Wireless Mesh Network Coverage Considerations	8-17
Cell Planning and Distance	8-17
For the Cisco 1520 Series Access Points	8-17
Collocating Mesh Access Points	8-17
Collocating AP1500s on Adjacent Channels	8-18
Collocating AP1500s on Alternate Adjacent Channels	8-18
CleanAir	8-18
CleanAir Advisor	8-19
Wireless Mesh Mobility Groups	8-19
Multiple Controllers	8-19
Increasing Mesh Availability	8-19
Multiple RAPs	8-21
Indoor Mesh Interoperability with Outdoor Mesh	8-22
Connecting the Cisco 1500 Series Mesh APs to the Network	8-22
Adding Mesh APs to the Mesh Network	8-23

CHAPTER 9

VoWLAN Design Recommendations 9-1

Antenna Considerations	9-1
AP Antenna Selection	9-1
Antenna Orientation	9-2
General Recommendations	9-4
Antenna Positioning	9-5
Handset Antennas	9-6
Channel Utilization	9-6
Dynamic Frequency Selection and 802.11h Requirements of the APs	9-7
5 GHz Band Channels	9-7
Call Capacity	9-9

AP Call Capacity	9-12
Cell Edge Design	9-14
Dual Band Coverage Cells	9-17
Dynamic Transmit Power Control	9-17
802.11r and 802.11k Features	9-18
Interference Sources Local to the User	9-19

CHAPTER 10

Cisco Unified Wireless Network Guest Access Services 10-1

Introduction	10-1
Scope	10-2
Wireless Guest Access Overview	10-2
Guest Access using the Cisco Unified Wireless Network Solution	10-2
WLAN Controller Guest Access	10-3
Supported Platforms	10-4
Auto Anchor Mobility to Support Wireless Guest Access	10-4
Anchor Controller Deployment Guidelines	10-5
Anchor Controller Positioning	10-5
DHCP Services	10-6
Routing	10-6
Anchor Controller Sizing and Scaling	10-6
Anchor Controller Redundancy	10-7
Web Portal Authentication	10-8
User Redirection	10-8
Guest Credentials Management	10-9
Local Controller Lobby Admin Access	10-10
Guest User Authentication	10-10
External Authentication	10-11
Guest Pass-through	10-11
Guest Access Configuration	10-12
Anchor WLC Installation and Interface Configuration	10-13
Guest VLAN Interface Configuration	10-14
Mobility Group Configuration	10-16
Defining the Default Mobility Domain Name for the Anchor WLC	10-16
Defining Mobility Group Members of the Anchor WLC	10-17
Adding the Anchor WLC as a Mobility Group Member of a Foreign WLC	10-18
Guest WLAN Configuration	10-18
Foreign WLC—Guest WLAN Configuration	10-20
Guest WLAN Configuration on the Anchor WLC	10-26
Anchor WLC—Guest WLAN Interface	10-26

Guest Account Management	10-28
Guest Management Using the Management System	10-28
Using the Add Guest User Template	10-30
Using the Schedule Guest User Template	10-33
Managing Guest Credentials Directly on the Anchor Controller	10-38
Configuring the Maximum Number of User Accounts	10-39
Maximum Concurrent User Logins	10-40
Guest User Management Caveats	10-41
Other Features and Solution Options	10-41
Web Portal Page Configuration and Management	10-41
Internal Web Page Management	10-42
Internal Web Certificate Management	10-44
Support for External Web Redirection	10-46
Anchor WLC-Pre-Authentication ACL	10-46
Anchor Controller DHCP Configuration	10-48
Adding a New DHCP Scope to the Anchor Controller	10-48
External Radius Authentication	10-50
Adding a RADIUS Server	10-50
External Access Control	10-54
Verifying Guest Access Functionality	10-55
Troubleshooting Guest Access	10-56
User Cannot Associate to the Guest WLAN	10-56
User Does Not Obtain an IP Address via DHCP	10-56
User is Not Redirected to Web Auth Page	10-57
User Cannot Authenticate	10-57
User Cannot Connect to Internet or Upstream Service	10-57
System Monitoring	10-57
Anchor Controller	10-57
Campus (Foreign) Controller	10-59
Debug Commands	10-61

CHAPTER 11

Cisco Mobility Services Engine 11-1

Introduction	11-1
Background Information	11-1
Overview	11-2
Terminology	11-2
Mobility Services Engine	11-2
Technology Background Information	11-4
RSSI	11-4

Time Difference on Arrival	11-5
Active RFID Tags	11-5
System Architecture	11-6
Related Information	11-8



Preface

Document Purpose

The purpose of this document is to describe the design and implementation of the Cisco Unified Wireless Network solution for the enterprise, using the features incorporated in the Cisco Wireless LAN Controller software Release 7.3.

Intended Audience

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

Document Organization

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

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

Section	Description
Chapter 6, “Cisco Unified Wireless Multicast Design.”	Describes the improvements that have been made in IP multicast forwarding and provides information on how to deploy multicast in a wireless environment.
Chapter 7, “FlexConnect.”	Describes the Cisco Centralized WLAN architecture and its use of H-REAP.
Chapter 8, “Cisco Wireless Mesh Networking.”	Describes the use of wireless mesh.
Chapter 9, “VoWLAN Design Recommendations.”	Provide design considerations when deploying voice over WLAN (VoWLAN) solutions.
Chapter 10, “Cisco Unified Wireless Network Guest Access Services.”	Describes the use of guest access services in the centralized WLAN architecture.
Chapter 11, “Cisco Mobility Services Engine.”	Discusses the Cisco Mobility Services Engine (MSE) and the areas that merit special consideration involving design, configuration, installation, and deployment.
Glossary	Lists and defines key terms used in the guide.



Cisco Unified Wireless Network Solution Overview

This chapter summarizes the benefits and characteristics of the Cisco Unified Wireless Network for the enterprise. The Cisco Unified Wireless Network solution offers secure, scalable, cost-effective wireless LANs for business critical mobility. The Cisco Unified Wireless Network is the industry's only unified wired and wireless solution to cost-effectively address the wireless LAN (WLAN) security, deployment, management, and control issues facing enterprises. This powerful indoor and outdoor solution combines the best elements of wired and wireless networking to deliver high performance, manageable, and secure WLANs with a low total cost of ownership.

WLAN Introduction

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

WLAN Solution Benefits

Benefits achieved by WLANs include:

- *Mobility within buildings or campus*—Facilitates implementation of applications that require an always-on network and that tend to involve movement within a campus environment.
- *Convenience*—Simplifies networking of large, open people-areas.
- *Flexibility*—Allows work to be done at the most appropriate or convenient place rather than where a cable drop terminates. Getting the work done is what is important, not where you are.
- *Easier to set-up temporary spaces*—Promotes quick network setup of meeting rooms, war rooms, or brainstorming rooms tailored to variations in the number of participants.
- *Lower cabling costs*—Reduces the requirement for contingency cable plant installation because the WLAN can be employed to fill the gaps.

- *Easier adds, moves, and changes and lower support and maintenance costs*—Temporary networks become much easier to set up, easing migration issues and costly last-minute fixes.
- *Improved efficiency*—Studies show WLAN users are connected to the network 15 percent longer per day than hard-wired users.
- *Productivity gains*—Promotes easier access to network connectivity, resulting in better use of business productivity tools. Productivity studies show a 22 percent increase for WLAN users.
- *Easier to collaborate*—Facilitates access to collaboration tools from any location, such as meeting rooms; files can be shared on the spot and requests for information handled immediately.
- *More efficient use of office space*—Allows greater flexibility for accommodating groups, such as large team meetings.
- *Reduced errors*—Data can be directly entered into systems as it is being collected, rather than when network access is available.
- *Improved efficiency, performance, and security for enterprise partners and guests*—Promoted by implementing guest access networks.
- *Improved business resilience*—Increased mobility of the workforce allows rapid redeployment to other locations with WLANs.

Requirements of WLAN Systems

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

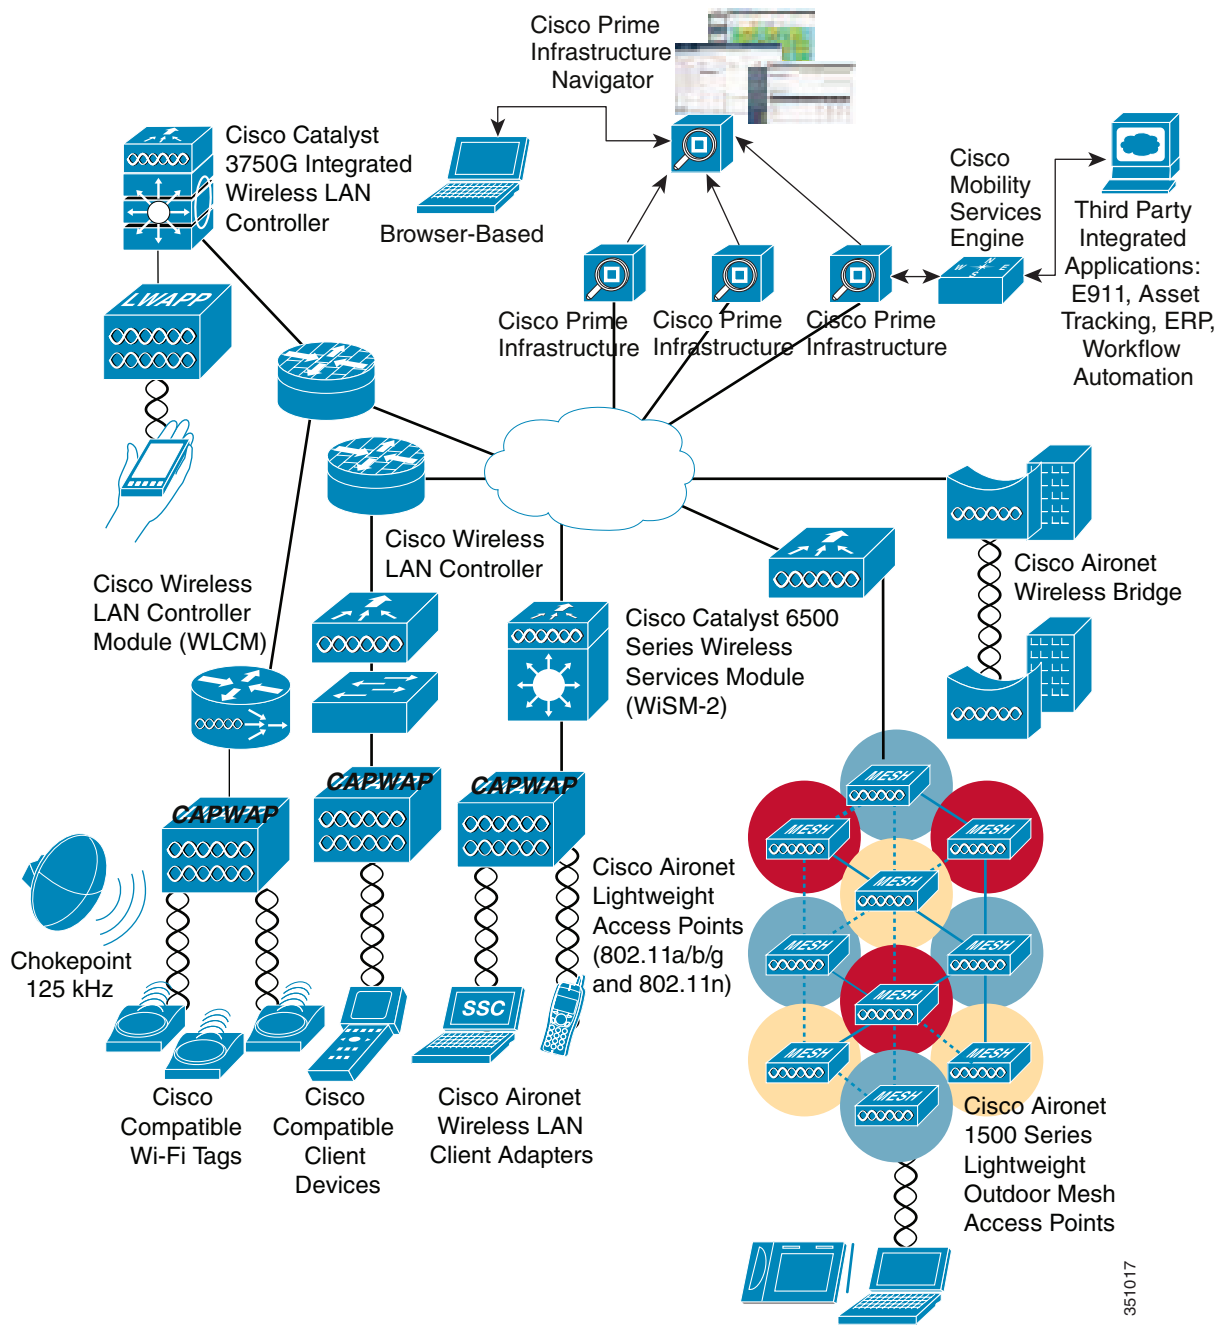
WLANs must be able to do the following:

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

- *Easily manage central or remote APs*—Network managers must be able to easily deploy, operate, and manage hundreds to thousands of APs within the WLAN campus deployments and branch offices or retail, manufacturing, and health care locations. The desired result is one framework that provides medium-sized to large organizations the same level of security, scalability, reliability, ease of deployment, and management that they have come to expect from their wired LANs.
- *Enhanced Security Services*—WLAN Intrusion Prevention System (IPS) and Intrusion Detection System (IDS) control to contain wireless threats, enforce security policy compliance, and safeguard information.
- *Voice Services*—Brings the mobility and flexibility of wireless networking to voice communications via the Cisco Unified Wired and Wireless network and the Cisco Compatible Extensions voice-enabled client devices.
- *Location Services* — Simultaneous tracking of hundreds to thousands of Wi-Fi and active RFID devices from directly within the WLAN infrastructure for critical applications such as high-value asset tracking, IT management, location-based security, and business policy enforcement.
- *Guest Access*— Provides customers, vendors, and partners with easy access to a wired and wireless LANs, helps increase productivity, facilitates real-time collaboration, keeps the company competitive, and maintains full WLAN security.

WLANs in the enterprise have emerged as one of the most effective means for connecting to a larger corporate network or to the internet. [Figure 1-1](#) shows the elements of the Cisco Unified Wireless Network.

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



The interconnected elements that work together to deliver a unified enterprise-class wireless solution include:

- Client devices
- Access points (APs)
- Network unification through controllers
- World-class network management

- Mobility services

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

Cisco Unified Wireless Network

The core components of Cisco Unified Wireless Networks include the:

- Aironet access points (APs)
- Wireless LAN controller (WLC)
- Cisco Prime Infrastructure
- Mobility Services Engine (MSE)

For more information about the Cisco Unified Wireless Network, see:

<http://www.cisco.com/go/unifiedwireless>

For information on optional Cisco components that provide additional benefits including *advanced enterprise-class security*, *extended RF management*, and *enhanced interoperability* see:

- Cisco Compatible Extensions client devices:
http://www.cisco.com/web/partners/pr46/pr147/partners_pgm_concept_home.html
- Cisco Secure Services clients:
<http://www.cisco.com/en/US/products/ps7034/index.html>



Cisco Unified Wireless Technology and Architecture

This chapter discusses the key design and operational considerations associated with the deployment of enterprise Cisco Unified Wireless Networks.

This chapter discusses:

- CAPWAP
- Core components
- Functional grouping of components
- Roaming
- Broadcast and multicast handling
- Design considerations
- Operation and maintenance

Much of the material in this chapter is explained in more detail in later chapters of this design guide. For more information on Cisco Unified Wireless technology, see the Cisco white paper on deployment strategies related to the Cisco 5500 Series Wireless LAN Controller at:

http://www.cisco.com/en/US/products/ps10315/prod_white_papers_list.html

CAPWAP Overview

Control and Provisioning of Wireless Access Points (CAPWAP) is the underlying protocol used in the Cisco *Centralized WLAN Architecture* (functional architecture of the Cisco Unified Wireless Network solution). CAPWAP provides for the configuration and management of WLANs, in addition to managing two-way tunneling traffic with WLAN clients to a centralized WLAN controller (WLC).

[Figure 2-1](#) shows a high-level diagram of a basic centralized WLAN deployment, where CAPWAP APs connect to a WLC by way of CAPWAP.

You can deploy CAPWAP controllers and LWAPP controllers on the same network. The CAPWAP-enabled software allows APs to join either a controller that runs CAPWAP or LWAPP. The only exception is the Cisco Aironet 1140 Series AP, which supports only CAPWAP and therefore joins only controllers that run CAPWAP. For example, an 1130 series AP can join a controller that runs either CAPWAP or LWAPP whereas an Aironet 1140 Series AP can join only a controller that runs CAPWAP.

Cisco recommends the following guidelines when using CAPWAP:

- If your firewall is currently configured to allow traffic only from APs that use LWAPP, you must change the rules of the firewall to allow traffic from APs that use CAPWAP.
- Ensure that the CAPWAP UDP ports 5246 and 5247 (similar to the LWAPP UDP ports 12222 and 12223) are enabled and are not blocked by an intermediate device that could prevent an AP from joining the controller.
- If access control lists (ACLs) are in the control path between the controller and its APs, you need to open new protocol ports to prevent APs from being stranded.

CAPWAP APs use a random UDP source port to reach these destination ports on the controller. In Cisco WLC release 5.2, LWAPP was removed and replaced by CAPWAP, but if you have a new out-of-the-box AP, it could try to use LWAPP to contact the controller before it downloads the CAPWAP image from the controller. Once the AP downloads the CAPWAP image from the controller, it uses only CAPWAP to communicate with the controller.

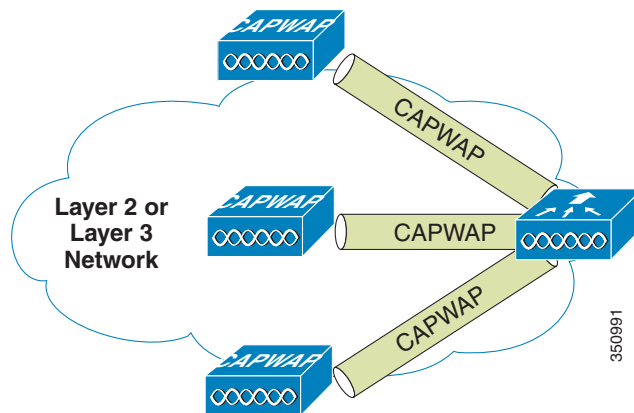


Note

After 60 seconds of trying to join a CAPWAP controller, the AP falls back to using LWAPP. If it cannot find a controller using LWAPP within 60 seconds, it tries again to join a controller using CAPWAP. The AP repeats this cycle of switching from CAPWAP to LWAPP and back again every 60 seconds until it joins a controller.

An AP with the LWAPP recovery image (an AP converted from autonomous (standalone) mode or an out-of-the-box AP) uses only LWAPP to try to join a controller before it downloads the CAPWAP image from the controller.

Figure 2-1 CAPWAP APs Connected to a WLC



Note

Although CAPWAP is made up of a number of functional components, only those that influence the design and operation of a centralized WLAN network are discussed in this design guide.

The key features of CAPWAP include:

- Split MAC tunnel
- L3-based tunnels
- WLC discovery process

Split MAC


A key component of CAPWAP is the concept of a *split MAC*, where part of the 802.11 protocol operation is managed by the CAPWAP AP, while the remaining parts are managed by the WLC. [Figure 2-2](#) is a diagram showing the split MAC concept.

A generic 802.11 AP, at the simplest level as shown in [Figure 2-2\(A\)](#), is nothing more than an 802.11 MAC-layer radio that bridges WLAN clients to a wired network based on association to a Basic Service Set Identifier (BSSID). The 802.11 standard extends the single AP concept (above) to allow multiple APs to provide an extended service set (ESS), as shown in [Figure 2-2\(B\)](#), where multiple APs use the same ESS identifier (ESSID, commonly referred to as an SSID) to allow a WLAN client to connect to a common network by way of more than one AP.

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

**Note**

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

Figure 2-2 CAPWAP Split MAC Concept

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

For example, the CAPWAP AP handles:

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

Other functionality is handled by the WLC. MAC layer functions provided by the WLC include:

- 802.11 authentication
- 802.11 association and reassociation (mobility)

350992

- 802.11 frame translation and bridging
- 802.1X/EAP/RADIUS processing
- Termination of 802.11 traffic on a wired interface, except in the case of REAP and H-REAP configured APs, which are discussed later in this design guide

A CAPWAP tunnel supports two categories of traffic:

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

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

- SSID
- Operational state
- Authentication and security method
- QoS

Layer 3 Tunnels

Cisco recommends using the Layer 3 CAPWAP tunnel type. This method uses IP UDP packets to facilitate communication between the CAPWAP AP and the WLC. Layer 3 CAPWAP is able to perform fragmentation and reassembly of tunnel packets. This allows client traffic to make use of a full 1500 byte MTU and not have to adjust for any tunnel overhead.



Note

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

The figures below are of Layer 3 CAPWAP packet captures used to illustrate CAPWAP operation. The sample decodes were captured using a Wireshark packet analyzer.



Note

By default, Wireshark does not decode Cisco CAPWAP packets correctly. Correct this in the Wireshark configuration window by selecting the *SWAP Frame Control* option under the *Protocol Preferences* tab.

Figure 2-3 shows a decode of a CAPWAP control packet. This packet originates from the WLC using UDP source port 5246 (as do all CAPWAP control packets from the WLC). Control Type 12 represents a configuration command used to pass AP configuration information to the CAPWAP AP by the WLC. Control packet payloads are AES encrypted, using keys derived from the PKI authentication process that is performed when a CAPWAP AP first establishes a connection with the WLC.

Figure 2-3 CAPWAP Control Packet

```

4 Frame 456: 165 bytes on wire (1320 bits), 165 bytes captured (1320 bits) on interface 0
4 Ethernet II, Src: Cisco_a9:91:94 (00:3a:a9:a9:91:94), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
4 Internet Protocol Version 4, Src: 172.20.227.125 (172.20.227.125), Dst: 255.255.255.255 (255.255.255.255)
4 User Datagram Protocol, Src Port: 39195 (39195), Dst Port: capwap-control (5246)
  Source port: 39195 (39195)
  Destination port: capwap-control (5246)
  Length: 131
  Checksum: 0x0000 (none)
4 Control And Provisioning of Wireless Access Points
  Preamble
  Header
    Header Length: 4
    Radio ID: 0
    wireless Binding ID: IEEE 802.11 (1)
  Header flags
    Fragment ID: 0
    Fragment Offset: 0
    Reserved: 0
    MAC length: 6
    MAC address: Cisco_dc:5a:00 (34:a8:4e:dc:5a:00)
    Padding for 4 Byte Alignment: 00
  Control Header
    Message Type: 1
    Sequence Number: 0
    Message Element Length: 102
    Flags: 0

```

351031

Figure 2-4 shows a decode of a CAPWAP packet containing an 802.11 probe request. This packet originates from the CAPWAP AP to the WLC using UDP port 5246 (as do all CAPWAP encapsulated 802.11 frames). In this example, received signal strength indication (RSSI) and signal-to-noise ratio (SNR) values are also included in the CAPWAP packet to provide RF information to the WLC.

Figure 2-4 CAPWAP 802.11 Probe Request

```

4 Frame 668: 151 bytes on wire (1208 bits), 151 bytes captured (1208 bits) on interface 0
4 Ethernet II, Src: Cisco_42:57:5c (44:d3:ca:42:57:5c), Dst: Cisco_da:78:20 (64:d8:14:da:78:20)
4 Internet Protocol Version 4, Src: 172.20.227.123 (172.20.227.123), Dst: 172.20.227.99 (172.20.227.99)
4 User Datagram Protocol, Src Port: 9590 (9590), Dst Port: capwap-data (5247)
  Source port: 9590 (9590)
  Destination port: capwap-data (5247)
  Length: 117
  Checksum: 0x0000 (none)
4 Control And Provisioning of Wireless Access Points
  Preamble
  Header
    Header Length: 4
    Radio ID: 0
    wireless Binding ID: IEEE 802.11 (1)
  Header flags
    1... .... = Payload Type: Native frame format (see wireless Binding ID field)
    .0.. .... = Fragment: Don't Fragment
    ..0. .... = Last Fragment: More fragments follow
    ...1 .... = Wireless header: Wireless Specific Information is present
    .... 0... = Radio MAC header: No Radio MAC Address
    .... 0.. . = Keep-Alive: No Keep-Alive
    .... ..00 0 = Reserved: Not set
  Fragment ID: 0
  Fragment Offset: 0
  Reserved: 0
  wireless length: 4
  wireless data: 00000000
  wireless data ieee80211 Frame Info: 00000000
    wireless data ieee80211 RSSI (dBm): 0
    wireless data ieee80211 SNR (dB): 0
    wireless data ieee80211 Data Rate (Mbps): 0
  Padding for 4 Byte Alignment: 000000
4 IEEE 802.11 Probe Request, Flags: .....

```

351032

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

Figure 2-5 CAPWAP 802.11 Data Frame

```

Internet Protocol Version 4, Src: 172.20.227.100 (172.20.227.100), Dst: 172.20.227.125 (172.20.227.125)
User Datagram Protocol, Src Port: capwap-data (5247), Dst Port: 39195 (39195)
  Source port: capwap-data (5247)
  Destination port: 39195 (39195)
  Length: 42
  Checksum: 0x0000 (none)
Control And Provisioning of wireless Access Points
  Preamble
  Header
    Header Length: 2
    Radio ID: 1
    wireless Binding ID: IEEE 802.11 (1)
  Header flags
    1... .... = Payload Type: Native frame format (see wireless Binding ID field)
    .0.. .... = Fragment: Don't Fragment
    ..0. .... = Last Fragment: More fragments follow
    ...0 .... = Wireless header: No wireless specific information
    .... 0... = Radio MAC header: No Radio MAC Address
    .... .0.. = Keep-Alive: No Keep-Alive
    .... ..00 0 = Reserved: Not set
  Fragment ID: 0
  Fragment Offset: 0
  Reserved: 0
IEEE 802.11 Disassociate, Flags: .....
  Type/Subtype: Disassociate (0x0a)
  Frame Control: 0x00A0 (Swapped)
    .000 0000 0000 0000 = Duration: 0 microseconds
  Destination address: Apple_d1:22:39 (18:20:32:d1:22:39)
  Source address: Cisco_dc:5a:00 (34:a8:4e:dc:5a:00)
  BSS Id: Cisco_dc:5a:00 (34:a8:4e:dc:5a:00)
  Fragment number: 0
  Sequence number: 0

```

351003

WLC Discovery and Selection

This section describes the typical behavior of a Layer 3 CAPWAP AP upon being reset.

For a comprehensive description of the discovery/join process, see the *Cisco Wireless LAN Controller Configuration Guide* at:

http://www.cisco.com/en/US/docs/wireless/controller/7.3/configuration/guide/b_cg73.html

CAPWAP AP Reset

Upon Layer 3 CAPWAP AP reset, the following sequence of steps takes place:

- Step 1** The AP broadcasts a Layer 3 CAPWAP discovery message on the local IP subnet. Any WLC configured for Layer 3 CAPWAP mode that is connected to the same IP subnet will see the discovery message. Each of the WLCs receiving the CAPWAP discovery message will in turn reply with a unicast CAPWAP discovery response message to the AP.
- Step 2** The AP maintains previously learned WLC IP addresses locally in NVRAM. The AP sends a unicast CAPWAP discovery request to each of these WLC IP addresses. Any WLC receiving a CAPWAP discovery request responds by sending a CAPWAP discovery response to the AP. As stated earlier, WLC IP addresses can be learned by way of OTAP messages sent from existing APs already joined to WLCs.

The information stored in NVRAM also includes address information for any previously joined WLC that was a member of another mobility group (for more information, see [Mobility Groups, AP Groups, and RF Groups, page 2-12](#)).

- Step 3** DHCP servers can be programmed to return WLC IP addresses using vendor-specific DHCP options. When programmed, *Option 43* is used in a *DHCP offer* to *advertise* WLC addresses to CAPWAP APs. When an AP receives its IP address by way of DHCP, it checks for WLC IP address information in the Option 43 field of the DHCP *offer*. The AP sends a unicast CAPWAP discovery message to each WLC listed in the DHCP Option 43. WLCs receiving the CAPWAP discovery request messages unicast a CAPWAP discovery response to the AP.
- Step 4** Without the Option 43 information, the AP attempts to resolve the DNS name `CISCO-CAPWAP-CONTROLLER.localdomain`. If the AP is able to resolve this, it sends a unicast CAPWAP discovery message to each IP address returned in the DNS reply. As described above, each WLC that receives a CAPWAP discovery request message replies with a unicast CAPWAP discovery response to the AP.
- Step 5** If, after Steps 1 through 4 no CAPWAP discovery response is received, the AP resets and restarts the search algorithm.
-

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

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

Core Components

The core components that make up the Cisco Unified Wireless Network solution are the Cisco Prime Infrastructure formally known as wireless control system (WCS) and network control system (NCS), wireless LAN controllers (WLCs), CAPWAP APs, and the Cisco Mobility Services Engine (MSEs). This section describes product options for Cisco Prime Infrastructure, WLCs, and APs (for information on MSEs, see [Chapter 11, “Cisco Mobility Services Engine.”](#)).

Cisco Wireless LAN Controllers

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

The following summarizes the various Cisco WLCs and their features:

- **Cisco 2504 WLC**—The 2504 controller works in conjunction with Cisco APs and Cisco Prime Infrastructure to provide system-wide wireless LAN functions. As a component of the Cisco Unified Wireless Network, the 2504 controller provides real-time communication between wireless APs and other devices to deliver centralized security policies, guest access, Wireless Intrusion Prevention System (WIPS), context-aware (location), award-winning RF management, quality of services for mobility services such as voice and video, and OEAP support for the teleworker solution.

Cisco 2504 WLCs support up to 50 APs in increments of five APs with a minimum of five APs, making it a cost-effective solution for retail, enterprise branches, and small and medium-sized businesses. The 2504 WLC comes with four Gigabit Ethernet ports.

- Cisco 5508 WLC—The 5508 controller provides real-time communication between Cisco Aironet AP, the Cisco Prime Infrastructure, and the Cisco Mobility Services Engine. It delivers centralized security policies, wireless intrusion prevention system (IPS) capabilities, RF management, and quality of service (QoS). To help ensure an exceptional end-user experience on the wireless network, the Cisco 5500 Series provides a variety of capabilities:
 - Integrated CleanAir technology protects 802.11n performance by enabling a self-healing, self-optimizing wireless network.
 - Cisco ClientLink technology optimizes mixed-client network capacity by helping ensure that 802.11a/g and 802.11n clients operate at the best possible rates.
 - The Cisco Identity Services Engine provides a single, centralized point of management across wired and wireless networks. Enterprises can adapt to the exponential growth in mobile smartphones, tablets, and laptops by providing secure and relevant access for employees, guests, and contractors.
- Cisco Wireless Services Module 2 (WiSM-2)—A WLC module that is designed specifically for the Cisco Catalyst 6500 switch series. It supports up to 1000 APs per module. Depending on the 6500 platform, multiple WiSM-2s can be installed to offer significant scaling capabilities. The WiSM-2 appears as a single aggregated link interface on the 6500 that can be configured as a dot1 trunk to provide connection into the 6500 backplane. This is ideal for large buildings or campuses.
- Cisco Virtual Wireless Controller (vWLC)—The Virtual Wireless Controller provides real-time, centralized communications between Cisco Aironet Access Points, the Cisco Prime Infrastructure, and the Cisco Mobility Services Engine. Designed for organizations with virtualization initiatives, and for small to medium enterprise deployments, the Virtual Wireless Controller offers:
 - Centralized wireless network visibility and control for up to 200 branch locations
 - Ability for IT managers to configure, manage, and troubleshoot up to 200 access points and 3000 clients via FlexConnect
 - Secure guest access, rogue detection for Payment Card Industry (PCI) compliance, and in-branch (locally switched) Wi-Fi voice and video
 - Reliable connectivity with the Cisco FlexConnect solution for branch network
 - Protection of access points connected to remote controllers from branch WAN failures; wireless clients remain connected with access to local resources
- Cisco 7500 WLC—The 7500 Series Controller provides the visibility and control needed to manage thousands of wireless branches from a single location. The controller's features are as follows:
 - Provides a cost-effective solution that does not require a local controller at each branch location.
 - Consolidated, remote management allows scaled and consistent control over thousands of branches.
 - Delivers secure, centralized policy management of distributed guest and employee access.
 - Helps ensure business continuity in each local branch through resiliency over WAN failures.
 - Efficient networking with local switching of data traffic allows WAN optimization and QoS policies without requiring tunneling across the WAN.
- Cisco 8500 WLC—The 8500 Series Controller provides real-time, centralized communications between Cisco Aironet Access Points, the Cisco Prime Infrastructure, and the Cisco Mobility Services Engine. Designed for service provider and large campus deployments, the 8500 Series Controller offers:
 - The industry's largest scalability in a single rack-unit space (1RU); a centralized touch point for up to 6000 access points, 64,000 clients, and 6000 branch locations

- High speed with 10 Gigabit Ethernet connectivity support: Two 10 Gigabit Ethernet ports for redundancy
- High-Availability with sub-second access point stateful fail-over ensuring SSIDs are highly available and minimal impact to wireless clients
- High resiliency with redundant dual power supplies

Table 2-1 summarizes the available Cisco WLCs.

Table 2-1 Summary of Cisco WLCs

	Cisco 2500 Wireless LAN Controller	Cisco 5508 Wireless LAN Controller	Cisco Flex 7500 Wireless LAN Controller	Cisco 8500 Wireless LAN Controller	Cisco WLAN Controller Module for Cisco Integrated Services Router	Cisco Catalyst 6500 Series Wireless Services Module 2 (WSM-2)
Controller Type	Standalone	Standalone	Standalone	Standalone	Module	Module
Platform Integration	N/A	N/A	N/A	N/A	2900 and 3900 Series Integrated Services Routers	Series Switches
Number of Lightweight Access Points Supported	5, 15, 25 or 50	12, 25, 50, 100, 250 or 500	250, 300, 500, 1000, 2000 or 3000	300-6,000	25 and 50	1,000
Number of clients Supported	500	7000	30,000	64,000	1000	15,000
	Remote location, branch office or campus	Remote location, branch office or campus	Branch/Remote location from the corporate location through a WAN link	SP Wi-Fi and Large Enterprise Campus	Remote location, branch office, or small office	Large campus
Uplink Interfaces	Four 1-Gbps ports	Eight 1-Gbps ports	2 x 10 Gigabit Ethernet interfaces	2 x 10 Gigabit Ethernet interfaces	One 10-/1---Mpps port	Eight 1-Gbps ports

351049

Cisco Access Points

Within the Cisco Unified Wireless Network there are two categories of Cisco APs: autonomous and CAPWAP APs. This section describes the various models of CAPWAP APs that are available.



Note

Cisco 1500 series MESH APs are mentioned briefly below but this design guide does not address wireless MESH applications or MESH deployment guidelines. For information about the Cisco MESH solution see: *Cisco Mesh Networking Solution Deployment Guide*, <http://www.cisco.com/en/US/docs/wireless/technology/mesh/7.3/design/guide/Mesh.html>.

CAPWAP APs

Table 2-2 summarizes the available Cisco CAPWAP APs.

Table 2-2 Summary of CAPWAP APs

	3600 Series	3500 Series	2600 Series	1260 Series	1140 Series	1040 Series	600 Series
DataRate	450 Mbps	300 Mbps	450 Mbps	300 Mbps	300 Mbps	300 Mbps	300 Mbps
RadioDesign	4x4:3	2x3:2	3x4:3	2x3:2	2x3:2	2x2:2	2x2:2
CleanAir	☑	☑	☑				
ClientLink	ClientLink 2.0	☑	ClientLink 2.0	☑	☑		
BandSelect	☑	☑	☑	☑	☑	☑	
VideoStream	☑	☑		☑	☑	☑	
Rogue AP Detection	☑	☑	☑	☑	☑	☑	
AdaptiveWIPS	☑	☑	☑	☑	☑	☑	
OfficeExtend	☑		☑				☑
FlexConnect	☑	☑	☑	☑	☑	☑	☑*
WirelessMesh	☑*	☑		☑	☑	☑	
Future-proofModularity	☑						
DataUplink(Mbps)	10/100/1000	10/100/1000	10/100/1000	10/100/1000	10/100/1000	10/100/1000	10/100
Power	802.3af	802.3af	802.3af	802.3af	802.3af	802.3af	100 to 240 VAC, 50-60 Hz
TemperatureRange in Celsius	(f) -0 to 40° C (e) -20 to 55° C	(f) -0 to 40° C (e) -20 to 55° C	(f) -0 to 40° C (e) -20 to 55° C	-20 to 55° C	-0 to 40° C	-0 to 40° C	0 to 40° C
Wi-Fi Standards	802.11 a/b/g/n	802.11 a/b/g/n	802.11 a/b/g/n	802.11 a/b/g/n	802.11 a/b/g/n	802.11 a/b/g/n	802.11 a/b/g/n

351050

Cisco Prime Infrastructure

Cisco Prime Infrastructure provides a single integrated solution for comprehensive life cycle management of the wired/wireless access, campus, and branch networks, and rich visibility into end-user connectivity and application performance assurance issues. Cisco Prime Infrastructure accelerates the rollout of new services, secure access and management of mobile devices, making “Bring Your Own Device” (BYOD) a reality for corporate IT. Tightly coupling client awareness with application performance visibility and network control, Cisco Prime Infrastructure helps ensure uncompromising end-user quality of experience. Deep integration with the Cisco Identity Services Engine (ISE) further extends this visibility across security and policy-related problems, presenting a complete view of client access issues with a clear path to solving them.

Cisco Prime Infrastructure is organized into a life cycle workflow that includes the following high-level task areas:

- **Design**—The design phase focuses on the overall design of feature or device patterns or templates. The design area is where you create reusable design patterns such as configuration templates. Cisco Prime Infrastructure provides predefined templates, but you can also create your own. These patterns and templates are intended for use in the deployment phase of the life cycle.
- **Deploy**—The deployment phase focuses on deploying previously defined designs or templates into your network. The deploy area is where you specify how to deploy features, making use of the templates created in the design phase. The deploy phase allows you to push configurations defined in your templates to one or many devices.
- **Operate**—The Operate area is where you monitor your network on a daily basis, as well as perform other day-to-day or ad hoc operations relating to network device inventory and configuration management. The Operate tab contains dashboards, the Device Work Center, and the tools you need for day-to-day monitoring, troubleshooting, maintenance, and operations.

- **Report**—Cisco Prime Infrastructure also provides reports that you can use to monitor the system and network health as well as troubleshoot problems. The Cisco Prime Infrastructure Report Launchpad provides report access and scheduling for all types of reporting functions.
- **Administration**—The Administration area is where you specify system configuration settings, manage access control, and specify data collection settings.

Mobility Groups, AP Groups, and RF Groups

Within the Cisco Unified Wireless Network there are three important *group* concepts:

- Mobility groups
- AP groups
- RF groups


The following sections describe the purpose and application of these groups within the Cisco Unified Wireless Network.

Mobility Groups

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

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

Figure 2-6 WLC Mobility Group



Mobility Group Definition

Creating a mobility group is simple and well documented. However, there are a few important considerations to keep in mind:

- Up to 24 regular WLCs (Cisco 2500, 5508, WiSM-2, 7500, 8500, virtual WLC, WLCM2 series) can be in a single mobility group. You can configure up to 24 Wireless Services Module (WiSM-2) blades in one mobility group. Therefore, up to a maximum of 24000 APs are supported in a single mobility group. An enterprise can consist of more WLCs and APs, but they must be configured as members of another mobility group.
- With WLC release 5.1, up to 72 WLCs can be in a single mobility group with up to 72000 APs (1000 APs per WLC).
- The WLCs do not have to be of the same model/type to be a member of a mobility group. A group can be comprised of any combination of Cisco 2500 Series controller, Cisco Flex 7500, Cisco 5500 Series Controller, Virtual Controller, 8500 series, WiSM-2, Cisco Wireless Controller Software for SRE, or Cisco Wireless LAN Controller Module but they should all be running the same software version. Although mobility groups can function with software differences between devices, Cisco strongly recommends you use a common software version to ensure feature and functional parity across a unified wireless deployment.
- A mobility group requires all WLCs in the group to use the same virtual IP address.
- Each WLC must use the same *mobility domain name* (group name) and be defined as a peer in each others *Static Mobility Members* list.
- For a wireless client to roam seamlessly between mobility group members (WLCs), a given WLAN SSID and security configuration must be configured identically across all WLCs comprising the mobility group.

Mobility Group Application

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

Mobility Group—Exceptions

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

If the auto anchor feature is used, no more than 71 WLCs can be mapped to a designated anchor WLC. Foreign WLCs do not, by virtue of being connected to the auto anchor, establish mobility relationships between each other. The anchor WLC must have a *static mobility group member* entry defined for each foreign WLC where a static mobility tunnel is needed. The same is true for each foreign WLC where a static mobility tunnel is being configured; the anchor WLC must be defined as a *static mobility group member* in the foreign WLC.

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

For a discussion on mobility anchor configuration, see [Chapter 10, “Cisco Unified Wireless Network Guest Access Services.”](#)

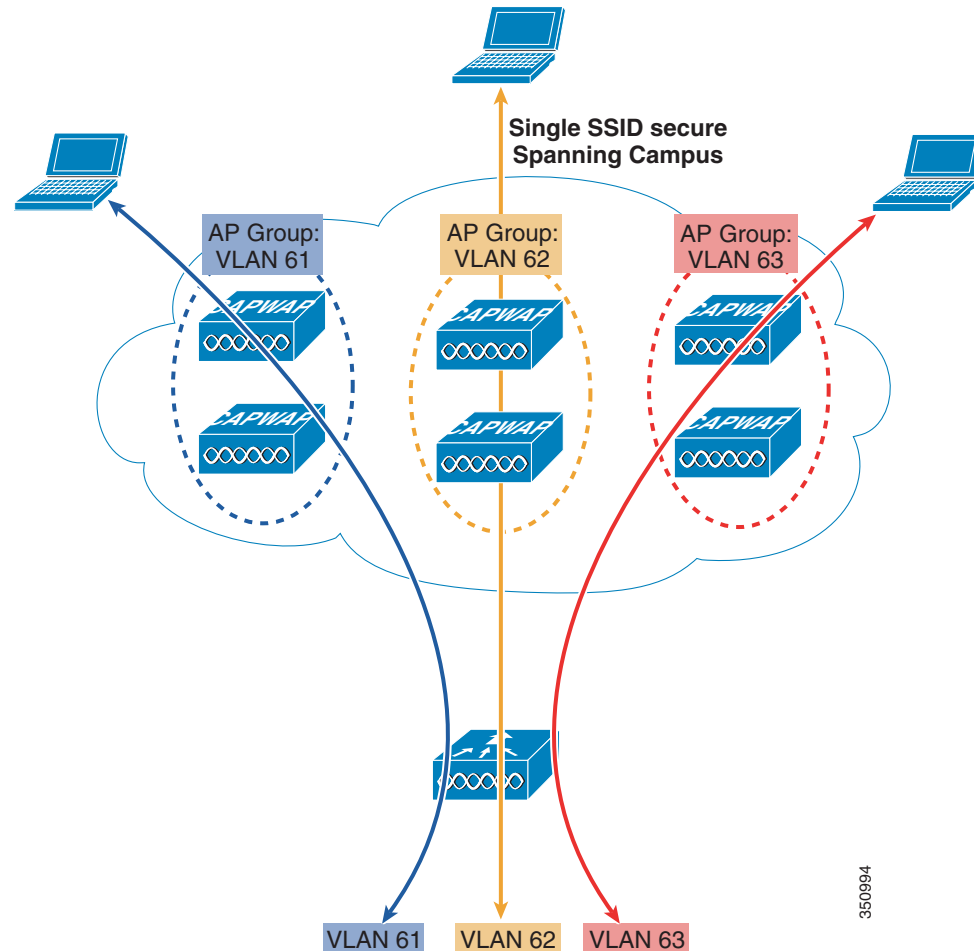
AP Groups

In typical deployment scenarios, each WLAN is mapped to a single dynamic interface per WLC. However, consider a deployment scenario where there is a 5508 WLC licensed to support maximum number of 500 APs. Now consider a scenario where 25 users are associated to each AP. That would result in 12,500 users sharing a single VLAN. Some customer designs might require substantially smaller subnet sizes. One way to deal with this is to break up the WLAN into multiple segments. The Cisco AP grouping feature allows a single WLAN to be supported across multiple dynamic interfaces (VLANs) on a WLC. This is done by taking a group of APs and mapping them to a specific dynamic interface. APs can be grouped logically by employee workgroup or physically by location. [Figure 2-7](#) illustrates the use of AP groups based on site-specific VLANs.



Note

AP groups do not allow multicast roaming across group boundaries. For more information, see [Chapter 6, “Cisco Unified Wireless Multicast Design”](#).

Figure 2-7 AP Groups and Site-Specific VLANs

As shown in [Figure 2-7](#), there are three dynamic interfaces configured, each mapping to a site-specific VLAN (VLAN 61, 62, and 63). Each site-specific VLAN and associated APs are mapped to the same WLAN SSID using the AP grouping feature. A corporate user associating to the WLAN on an AP in the AP group corresponding to VLAN 61 is assigned an IP address on the VLAN 61 IP subnet. Likewise, a corporate user associating to the WLAN on an AP in the AP group corresponding to VLAN 62 is assigned an IP address on the VLAN 62 IP subnet and so on. Roaming between the site-specific VLANs is handled internally by the WLC as a Layer 3 roaming event and because of this the wireless LAN client maintains its original IP address.

RF Groups

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


An RF group exists for each 802.11 PHY type. Grouping WLCs into RF domains allows the solution's dynamic RRM algorithms to scale beyond a single WLC, thereby allowing RRM for a given RF domain to extend between floors, buildings, and even across campuses. RF Groups and RRM is discussed in more detail in [Chapter 3, "WLAN RF Design Considerations,"](#) but can be summarized as follows:

- CAPWAP APs periodically send out neighbor messages over the air that includes the WLC IP address and a hashed message integrity check (MIC) derived from a timestamp and the BSSID of the AP.
- The hashing algorithm uses a shared secret (the RF Group Name) that is configured on the WLC and is pushed out to each AP. APs sharing the same secret are able to validate messages from each other using the MIC. When APs belonging to other WLCs hear validated neighbor messages at a signal strength of -80 dBm or stronger, their WLCs dynamically become members of the RF group.
- Members of an RF group elect an RF domain leader to maintain a *master* power and channel scheme for the RF group.
- The RF group leader analyzes real-time radio data collected by the system and calculates a master power and channel plan.
- The RRM algorithms attempt to:
 - Achieve a uniform (optimal) signal strength of -65 dBm across all APs
 - Avoid 802.11 co-channel interference and contention
 - Avoid non-802.11 interference.
- The RRM algorithms employ dampening calculations to minimize system-wide dynamic changes. The end result is dynamically calculated, near-optimal power and channel planning that is responsive to an ever changing RF environment.
- The RF group leader and members exchange RRM messages at a specified update interval, which is 600 seconds by default. Between update intervals the RF group leader sends keep alive messages to each of the RF group members and collects real-time RF data. Note that the maximum number of WLCs per RF group is 20.

Roaming

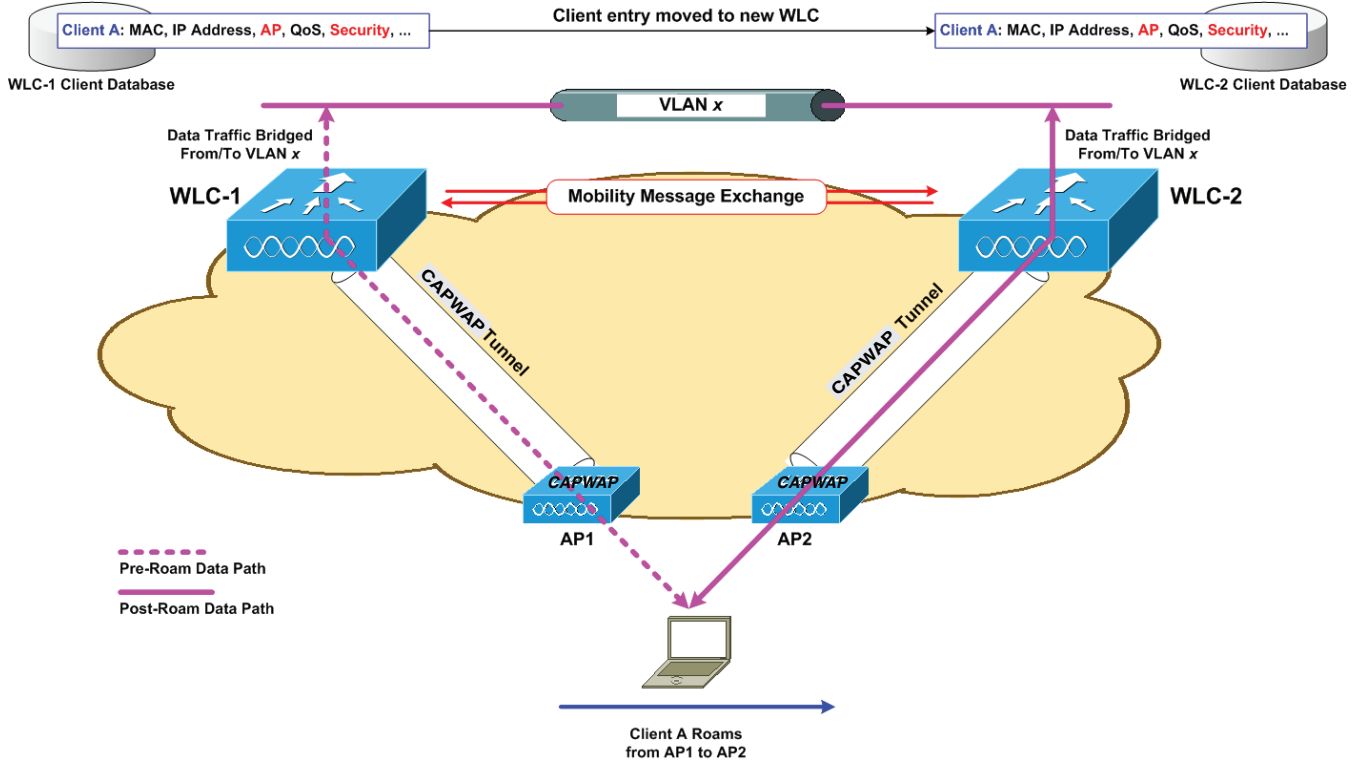
Mobility, or roaming, is the ability of a wireless LAN client to maintain its association seamlessly from one AP to another securely and with as little latency as possible. This section explains how mobility works when controllers are included in a wireless network.

When a wireless client associates and authenticates to an AP, the controller places an entry for that client in its client database. This entry includes the client MAC and IP addresses, security context and associations, quality of service (QoS) contexts, the WLAN, SSID and the associated AP. The controller uses this information to forward frames and manage traffic to and from the wireless client. [Figure 2-8](#) shows a wireless client that roams from one AP to another when both APs are joined to the same controller.

Figure 2-8 Intra-Controller Roaming

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

The process becomes more complicated, however, when a client roams from an AP joined to one controller to an AP joined to a different controller. It also varies based on whether the controllers are operating on the same subnet. Figure 2-9 shows inter-controller roaming, which occurs when the controllers' wireless LAN interfaces are on the same IP subnet.


Figure 2-9 Inter-Controller Roaming

When the client associates to an AP joined to a new controller, the new controller exchanges mobility messages with the original controller, and the client database entry is moved to the new controller. New security context and associations are established if necessary, and the client database entry is updated for the new AP. This process remains transparent to the user.

**Note**

All clients configured with 802.1X/Wi-Fi Protected Access (WPA) security complete a full authentication in order to comply with the IEEE standard.

Figure 2-10 shows inter-subnet roaming, which occurs when the controllers' wireless LAN interfaces are on different IP subnets.

Figure 2-10 Inter-Subnet Roaming

Inter-subnet roaming is similar to inter-controller roaming in that the controllers exchange mobility messages on the client roam. However, instead of moving the client database entry to the new controller, the original 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 controller. The roam remains transparent to the wireless client, and the client maintains its original IP address.

In inter-subnet roaming, WLANs on both anchor and foreign controllers need to have the same network access privileges and no source-based routing or source-based firewalls in place. Otherwise, the clients may have network connectivity issues after the handoff.

**Note**

If a client roams in web authentication state, the client is considered as a new client on another controller instead of considering it as a mobile client.

**Note**

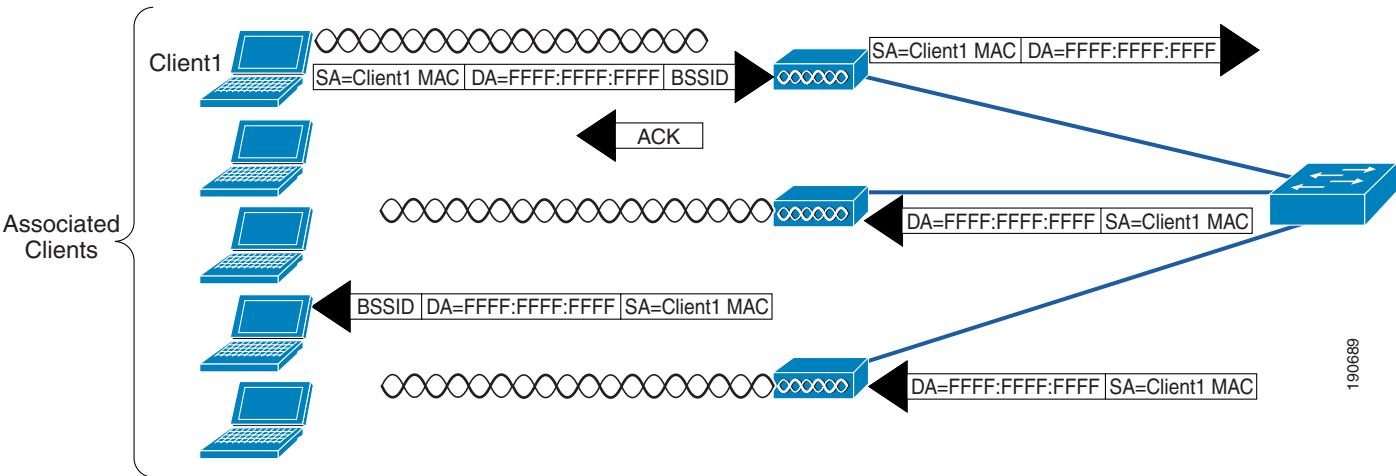
Seamless mobility is not supported for native IPv6 clients if the interface is untagged.

Broadcast and Multicast on the WLC

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

Figure 2-11 depicts basic 802.11 broadcast/multicast behavior. In this example, when Client 1 sends an 802.11 broadcast frame it is unicasted to the AP. The AP then sends the frame as a broadcast out both of its wireless and wired interfaces.

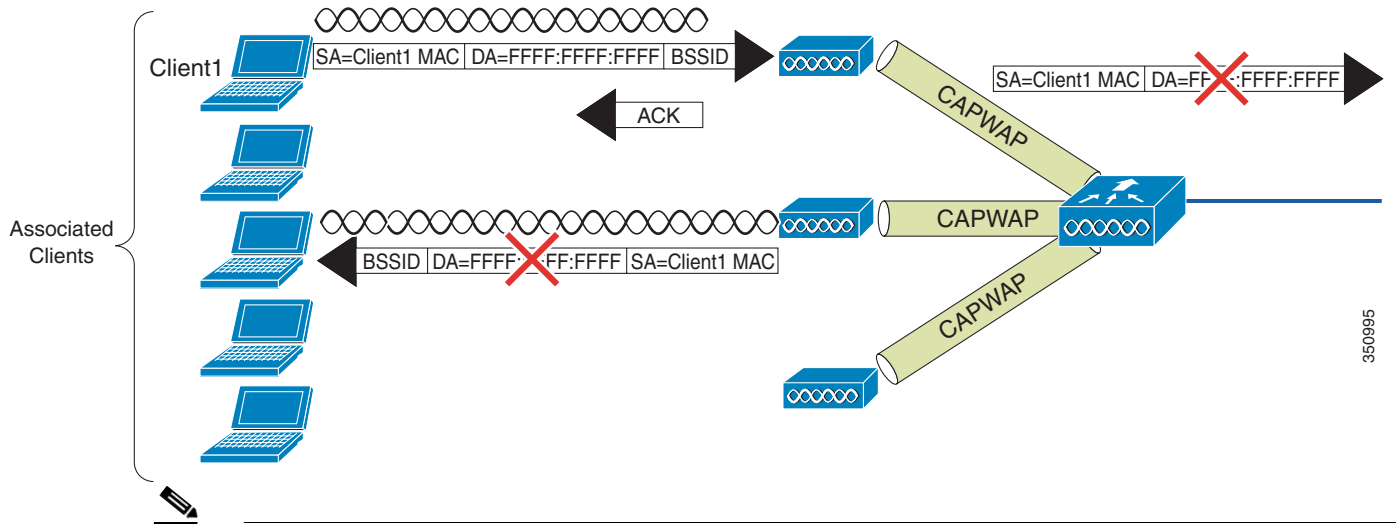
Figure 2-11802.11 Broadcast/Multicast



If there are other APs on the same wired VLAN as the AP, as shown in Figure 2-11, they forward the wired broadcast packet out their wireless interface.

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

Figure 2-12Default WLC Broadcast Behavior



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

WLC Broadcast and Multicast Details

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

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

DHCP

The WLC acts as a DHCP relay agent for associated WLAN clients. The WLC unicasts client DHCP requests to a locally configured or upstream DHCP server except during Layer 3 client roaming (discussed in more detail below). DHCP server definitions are configured for each dynamic interface, which in turn is associated with one or more WLANs. DHCP relay requests are forwarded by way of the dynamic interfaces using the source IP address of a given dynamic interface. Because the WLC knows which DHCP server to use for a given interface/WLAN, there is no need to broadcast client DHCP requests out its wired and wireless interfaces.

This method accomplishes the following:

- It eliminates the need for DHCP requests to be broadcasted beyond the WLC.
- The WLC becomes part of the DHCP process, thereby allowing it to learn the MAC/IP address relationships of connected WLAN clients, which in turn allows the WLC to enforce DHCP policies and mitigate against IP spoofing or denial-of-service (DoS) attacks.

VideoStream

The VideoStream feature makes the IP multicast stream delivery reliable over the air, by converting the broadcast frame over the air to a unicast frame. Each VideoStream client acknowledges receiving a video IP multicast stream. VideoStream is supported on all Cisco APs.

The following are the recommended guidelines for configuring VideoStream on the controller:

- The AP1100 and AP1200 do not support the reliable multicast feature.
- Ensure that the multicast feature is enabled. Cisco recommends configuring IP multicast on the controller with multicast-multicast mode.
- Check for the IP address on the client device. The device should have an IP address from the respective VLAN.
- Verify that the AP has joined the controllers.
- Ensure that the clients are able to associate to the configured WLAN at 802.11a/n speed.

Other Broadcast and Multicast Traffic

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

All normal precautions should be taken to limit the multicast address groups explicitly supported by a WLAN. When multicast is enabled, it is global in nature, meaning it is enabled for every WLAN configured regardless if multicast is needed by that WLAN or not. The Cisco Unified Wireless Network solution is not able to distinguish between data link layer and network layer multicast traffic, nor is the WLC capable of filtering specific multicast traffic. Therefore, the following additional steps should be considered:

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

Design Considerations

For Cisco Unified Wireless Network deployments, the primary design considerations are WLC location and AP and WLC connectivity. This section will briefly discuss these topics and make general recommendations where appropriate.

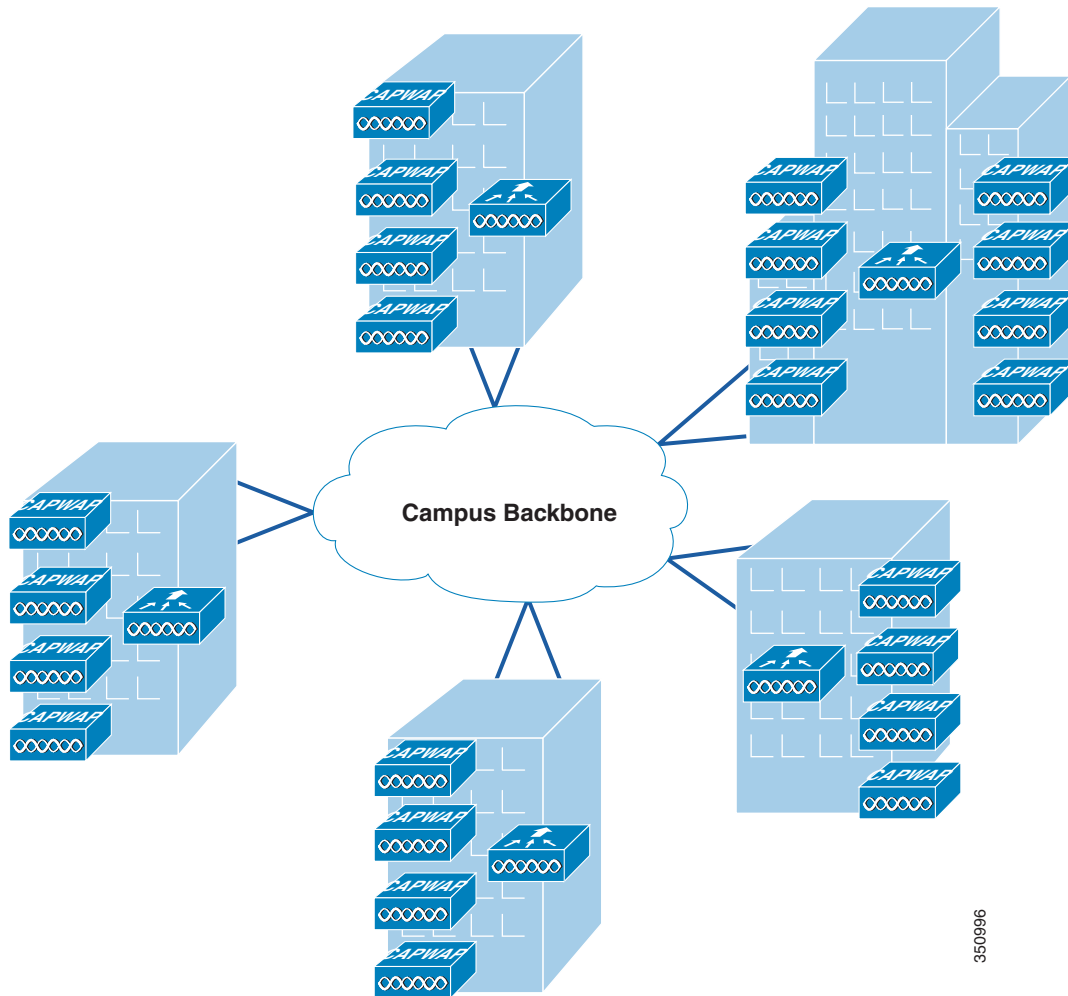
WLC Location

The Cisco Unified Wireless Network solution allows you to locate the WLCs in a distributed or centralized deployment, as described in the following sections.

Distributed WLC Deployment

[Figure 2-13](#) illustrates a distributed WLC deployment. In this model the WLCs are located throughout the campus network, typically on a per building basis, to manage the APs that are resident in the given building. The WLCs are connected to the campus network using the distribution routers within the building. In this scenario the CAPWAP tunnels, between APs and the WLC, typically stay within the building.

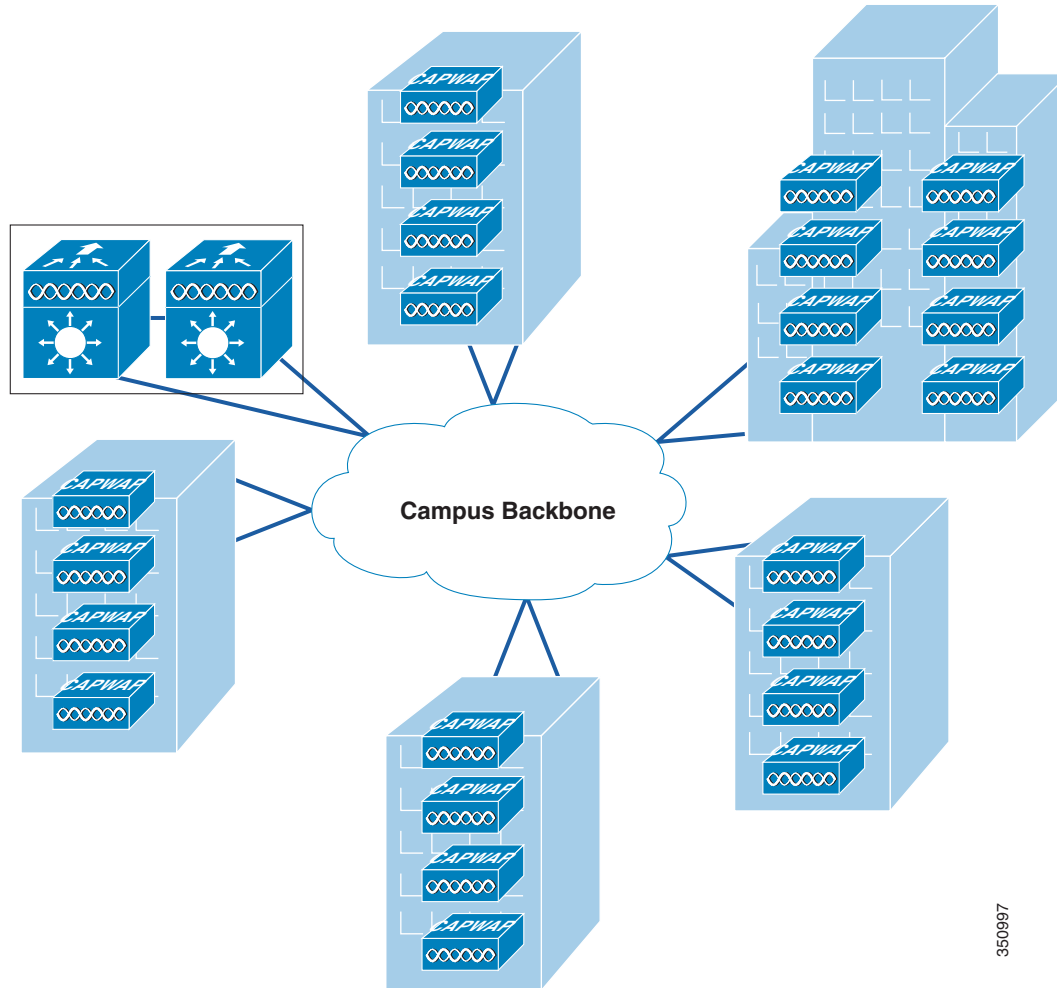
Each of the distributed WLCs could be configured as a separate RF group and Mobility group, so long as the WLAN coverage is not overlapping between buildings.

Figure 2-13 WLCs in Distributed Deployment

350996

Centralized WLC Deployment

Figure 2-14 illustrates a centralized WLC deployment. In this model, WLCs are placed at a centralized location in the enterprise network. This deployment model requires the AP/WLC CAPWAP tunnels to traverse the campus backbone network. Note in the illustration that the centralized WLCs are not shown in a specific building. A centralized WLC cluster is connected by way of a dedicated switch block to the campus core, which is typically located in the same building as the data center. The WLCs should not be connected directly to the data center switching block because the network and security requirements of a data center are generally different than that of a WLC cluster.

Figure 2-14 WLCs in Centralized Deployment

350997

Centralizing WLCs

Cisco generally recommends that you deploy the WLCs at a central location within the overall campus environment. The distributed deployment model (which would require mobility groups and Layer 3 roaming) is well proven, but it is not recommended because of current shortcomings with multicast support associated with Layer 3 roaming. When these are addressed, most of the barriers preventing consideration of a distributed deployment model will be removed.

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

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


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

In summary, a centralized WLC deployment is the preferred and recommended method. When planning any centralized WLC deployment, consideration must be given to the protection of the wired network infrastructure that directly connects to the WLC. This is because the WLC essentially attaches an *access network* at a location within the overall enterprise topology that would not otherwise be exposed to an *access network* and its associated vulnerabilities. Therefore, all security considerations normally associated with an access layer network device must be considered. For example, in a WISM-2-based deployment, features such as DoS protection and traffic storm protection should be considered because of the large-scale role the WISM-2 plays in providing diverse WLAN services to large numbers of end users while at the same time being directly connected to the backplane of a core multi-layer, multi-function Catalyst 6500 switching platform.

Distributed WLC Network Connectivity

A Layer 3 connected WLC, as shown in [Figure 2-15](#) (in this case a Cisco 3750G), allows the WLAN-related software and configuration to be isolated to a single device and connects to the distribution layer using the same routing configuration as access layer routing devices.

Figure 2-15 **Layer 3 Connected WLC**



Traffic Load and Wired Network Performance

When deploying a Cisco Unified Wireless Network solution, topics of discussion often include:

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

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

- Volume of CAPWAP control traffic
- Overhead introduced by tunneling
- Traffic engineering

Volume of CAPWAP Control Traffic

The volume of traffic associated with CAPWAP control can vary depending on the actual state of the network. For example, traffic volume is usually higher during a software upgrade or WLC reboot situations. Traffic studies have found that the average load CAPWAP control traffic places on the network is approximately 0.35 Kbps. In most campuses, this would be considered negligible and would be of no consequence when considering a centralized deployment model over a distributed one.

Overhead Introduced by Tunneling

A CAPWAP tunnel adds 44 bytes to a typical IP packet to and from a WLAN client. Given that average packets sizes found on typical enterprises are approximately 300 bytes, this represents an overhead of approximately 15 percent. In most campuses, this overhead would be considered negligible and again would be of no consequence when considering a centralized deployment model over a distributed one.

Traffic Engineering

Any WLAN traffic that is tunneled to a centralized WLC is then routed from the location of the WLC to its end destination in the network. Depending on the distance of the tunnel and location of the WLC, WLAN client traffic might not otherwise follow an optimal path to a given destination. In the case of a traditional access topology or distributed WLC deployment, client traffic enters the network at the edge and is optimally routed from that point based on destination address.

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

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

AP Connectivity

APs should be on different subnets from the end users (802.11 clients). This is consistent with general best-practice guidelines that specify that infrastructure management interfaces should be on a separate subnet from end users. Additionally, Cisco recommends that Catalyst Integrated Security Features (CISF) be enabled on the CAPWAP AP switch ports to provide additional protection to the WLAN infrastructure. (FlexConnect AP connectivity is discussed in [Chapter 7, “FlexConnect.”](#))

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

In order to effectively offer WLAN QoS within the Cisco Unified Wireless Network, QoS should also be enabled throughout the wired network that provides connectivity between CAPWAP APs and the WLCs.

Operation and Maintenance

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

WLC Discovery

The different WLC discovery mechanisms for APs (discussed earlier) make initial deployment of CAPWAP APs very simple. Options include:

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

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

AP Distribution

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

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

The elements used to control AP distribution across multiple WLCs are:

- Primary, secondary, and tertiary WLC names—Each AP can be configured with a primary, secondary, and tertiary WLC name, which in turn determines the first three WLCs in the mobility group that the AP will prefer to join regardless of the load variations across WLCs in the mobility group.

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



WLAN RF Design Considerations

This chapter describes the basic information necessary to understand radio frequency (RF) considerations in planning for various wireless local area network (WLAN) environments. The topics of this chapter include:

- Regulatory domains and RF considerations
- IEEE 802.11 standards
- RF spectrum implementations of 802.11b/g/n (2.4-GHz) and 802.11a/n/ac (5-GHz)
- Planning for RF deployment
- Manually fine-tuning WLAN coverage
- Radio resource management (RRM) algorithms

RF Basics

In the United States there are three bands (frequency ranges) allocated for unlicensed industrial, scientific, and medical (ISM) usage (see [Figure 3-1](#)).

The ISM bands are designated as the:

- 900 MHz band: 902 to 928 MHz
- 2.4 GHz band (IEEE 802.11b/g/n): 2.4 to 2.4835 GHz
- 5 GHz band (IEEE 802.11a/n/ac):
 - 5.150 to 5.250 GHz (UNII-1)
 - 5.250 to 5.350 GHz (UNII-2)
 - 5.725 to 5.875 GHz (UNII-3/ISM)

Each band has different characteristics. The lower frequency 2.4 GHz band exhibits better range but with limited bandwidth and thus lower data rates. The higher frequency 5 GHz band exhibits less range and is subject to greater attenuation from solid objects but has higher data rates.

The following sections provide a summary of regulatory domains and their operating frequencies.

Regulatory Domains

Devices that operate in unlicensed bands do not require a formal licensing process, but when operating in the ISM bands, the vendor is obligated to follow the government regulations for that region. The regulatory agencies in different parts of the world monitor these bands according to different criteria. WLAN devices must comply with the specifications of the relevant governing regulatory body. Although the regulatory requirements do not affect the interoperability of IEEE 802.11b/g/n- and 802.11a/n/ac-compliant products, the regulatory agencies do set certain criteria in the standard. For example, the emission requirements for WLAN are to minimize the amount of interference a radio can generate or receive from another radio in the same proximity. It is the responsibility of the WLAN vendor to obtain product certification from the relevant regulatory body.

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

Operating Frequencies

The 2.4-GHz band regulations of 802.11b/g/n have been relatively constant, given the length of time they have been in operation. The FCC (U.S) allows for 11 channels, ETSI (most other parts of the world) allows for up to 13 channels, and Japan allows up to 14 channels but requires a special license to operate in channel 14. Countries that adhere to the 5.0-GHz band regulations of 802.11a/n/ac are moving to open additional spectrum for additional channels.

These frequency bands and their associated data rates are described in more detail in the following sections.

2.4 GHz - 802.11b/g/n

The 802.11b standard, which was ratified in 1999, supports data rates of 1, 2, 5.5, and 11 Mbps and enjoys broad user acceptance and vendor support. 802.11b has been deployed by thousands of enterprise organizations as the first standardized method of WLAN communication.

The 802.11g standard, which was ratified in 2003, operates in the same spectrum as and is backwardly compatible with 802.11b. The 802.11g standard supports the additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps and is the most popular method of WLAN communications in the 2.4 GHz band.

The 802.11n standard, which was ratified in 2009, allows for usage in both 2.4 and 5 GHz bands. The 2.4 GHz band supports data rates up to 144 Mbps (assuming 20 MHz bandwidth and single transmitter stream). Note that faster speeds and bonding channels (using spectrum greater than 20 MHz) can permit speeds up to 300 Mbps, but this is typical of a home deployment. Although enterprise networks are limited to a data rate of 20 MHz in the 2.4 GHz band (due to a limited amount of spectrum), faster speeds can be achieved using the 5 GHz band with 802.11n and/or 802.11ac technology.

5 GHz - 802.11a/n/ac

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

Deployment Considerations

Choosing between these two technologies (802.11a/n/ac and 802.11b/g/n) does not involve an equal one-for-one trade-off. They are complementary technologies and will continue to coexist in future enterprise environments. Those responsible for implementing these technologies must be able to make an educated choice between deploying 2.4 GHz-only networks, 5 GHz-only networks, or a combination of them. Organizations with existing 802.11b/g networks cannot simply deploy a new 802.11a network for existing APs and expect to have their 54 Mbps 802.11a/n/ac coverage in the same areas as their 11 Mbps 802.11b/g/n coverage. The technical characteristics of both these bands simply do not allow for this kind of interchangeability in coverage.


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

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

There are differing limitations on these three UNII bands with restrictions varying between them for transmit power, antenna gain, antenna styles, and usage. The UNII-1 band is designated for indoor operations and initially had a restriction of permanently attached antennas. The UNII-2 and UNII-3 bands are designated for indoor or outdoor operations and permit external antennas.

The channels in UNII-1 (5.150 to 5.250 GHz) are 36, 40, 44, and 48. The channels in UNII-2 (5.250 to 5.350 GHz) are 52, 56, 60, 64 and require dynamic frequency selection (DFS) and transmitter power control (TPC). The channels in the new frequency range (5.470 to 5.725 GHz) are 100, 104, 108, 112, 116, 120, 124, 128, 132, 136, and 140 also require DFS and TPC. The channels in UNII-3 (5.725 to 5.850) are 149, 153, 157, 161, and 165 and do not require DFS and TPC. Not all channels in a given range can be used in all of the regulatory domains. [Figure 3-1](#) shows the various channels in the UNII-1, 2, and 3 bands, along with the additional 11 new channels.

Figure 3-1 802.11 Channel Capacity



The FCC released a revision to the regulations in 2004 covering the 5 GHz 802.11a channel usage. This revision added 11 additional channels, bringing the available channels capacity to 23 channels, as shown in [Figure 3-1](#). The 11 new channels are for indoor/outdoor use. However, to use the new channels radios must comply with the DFS and TPC features that are part of the 802.11h specification. DFS is required

to avoid radar that operates in this frequency range, but it can also be used for other purposes such as dynamic frequency planning. 802.11h has been supported by Cisco Unified Wireless Network since 2010.

DFS dynamically instructs a transmitter to switch to another channel whenever a particular condition such as the presence of a radar signal is met. Before transmitting, the DFS mechanism of a device monitors its available operating spectrum and listens for a radar signal. If a signal is detected the channel associated with the radar signal is vacated or flagged as unavailable for use by the transmitter. Prior to and during operation the transmitting device continuously monitors the environment for the presence of radar. Portions of the 5 GHz band are allocated to radar systems, which allows WLANs to avoid interference with incumbent radar users in instances where they are co-located.

TPC allows the AP to negotiate power levels with a WLAN client during the association process. The AP can inform the WLAN client of the range of allowable transmit power to be used with that AP and might reject clients unable to meet those levels. The WLAN client is able to adjust its transmit power level within the range specified in the TPC negotiations. This ensures that interference from the WLAN is minimized and allows the WLAN client to optimize battery life.

For more information on FCC regulation updates, see the Cisco white paper *FCC Regulations Update* at: http://www.cisco.com/en/US/products/hw/wireless/ps469/products_white_paper0900aecd801c4a88.shtml

Table 3-1 shows the standard 802.11a/n/ac frequencies.

Table 3-1 Operating Frequency Range for 802.11a/n/ac

Channel ID	36	40	44	48	52	56	60	64	100	104	108	112	116	120	124	128	132	136	140	149	153	157	161	165	
Center Freq. in MHz	5180	5200	5220	5240	5260	5280	5300	5320	5500	5520	5540	5560	5580	5600	5620	5640	5660	5680	5700	5745	5765	5785	5805	5825	
Band	UNII-1				UNII-2																	UNII-3			ISM

Understanding the IEEE 802.11 Standards

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

For more information on these working groups see: <http://www.ieee802.org/11/>

Table 3-2 IEEE 802.11 Task Group Activities

Task Group	Project
MAC	Develop one common MAC for WLANs in conjunction with a physical layer entity (PHY) task group.
PHY	Develop three WLAN PHYs—Infrared, 2.4 GHz FHSS, 2.4 GHz DSSS.
a	Develop PHY for 5 GHz UNII band.
b	Develop higher rate PHY in 2.4 GHz band.
c	Cover bridge operation with 802.11 MACs (spanning tree).
d	Define physical layer requirements for 802.11 operation in other regulatory domains (countries).

Table 3-2 IEEE 802.11 Task Group Activities (continued)

e	Enhance 802.11 MAC for QoS.
f	Develop recommended practices for Inter Access Point Protocol (IAPP) for multi-vendor use.
g	Develop higher speed PHY extension to 802.11b (54 Mbps).
h	Enhance 802.11 MAC and 802.11a/n/ac PHY-Dynamic Frequency selection (DFS), Transmit Power control (TPC).
i	Enhance 802.11 MAC security and authentication mechanisms.
j	Enhance the 802.11 standard and amendments to add channel selection for 4.9 GHz and 5 GHz in Japan.
k	Define RRM enhancements to provide interfaces to higher layers for radio and network measurements.
k	Define Radio Resource Measurement enhancements to provide interfaces to higher layers for radio and network measurements.
m	Perform editorial maintenance, corrections, improvements, clarifications, and interpretations relevant to documentation for 802.11 family specifications.
n	Focus on high throughput extensions (>100 Mbps at MAC SAP) in 2.4 GHz and/or 5 GHz bands.
o	Provide Fast Handoffs in Voice over WLAN (goal is around 50 ms)
p	Focus on vehicular communications protocol aimed at vehicles, such as toll collection, vehicle safety services, and commerce transactions via cars.
r	Develop a standard specifying fast BSS transitions and fast roaming.
s	Define a MAC and PHY for meshed networks that improves coverage with no single point of failure.
t	Provide a set of performance metrics, measurement methodologies, and test conditions to enable manufacturers, test labs, service providers, and users to measure the performance of 802.11 WLAN devices and networks at the component and application level.
u	Provide functionality and interface between an IEEE 802.11 access network (Hotspot) and any external network.
v	Provide extensions to the 802.11 MAC/PHY to provide network management for stations (STAs).
w	Provide mechanisms that enable data integrity, data origin authenticity, replay protection, and data confidentiality for selected IEEE 802.11 management frames including but not limited to: action management frames, deauthentication and disassociation frames.

Direct Sequence Spread Spectrum (DSSS)


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

IEEE 802.11b Direct Sequence (DS) Channels

Fourteen channels are defined in the IEEE 802.11b direct sequence (DS) channel set. Each DS channel transmitted is 22 MHz wide but the channel separation is only 5 MHz. This leads to channel overlap so that signals from neighboring channels that are less than are 25 MHz apart can interfere with each other. In the 14-channel DS system (11 usable channels in the US), the only non-overlapping, non-interfering channels possible are channels 1, 6, and 11 (see [Figure 3-2](#)).

This channel spacing governs the use and allocation of channels in a multi-AP environment such as an office or campus. APs are usually deployed in a cellular fashion within an enterprise where adjacent APs are allocated non-overlapping channels. Alternatively, APs can be co-located using channels 1, 6, and 11 to deliver 33 Mbps bandwidth to a single area (but only 11 Mbps to a single client). If 802.11g is used in the same manner the aggregate bandwidth is 162 Mbps with a maximum data rate of 54 Mbps. This channel allocation scheme is illustrated in [Figure 3-2](#).

Figure 3-2 IEEE 802.11 DSS Channel Allocations



IEEE 802.11g

802.11g provides for a higher data rate (up to 54 Mbps) in the 2.4 GHz band, the same spectrum as 802.11b. 802.11g is backward-compatible with 802.11b and provides additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. At higher data rates, 802.11g uses the same modulation technique, orthogonal frequency division multiplexing (OFDM), as 802.11a/n/ac (see [IEEE 802.11a OFDM Physical Layer](#), page 3-7).

[Table 3-3](#) lists 802.11g modulation and transmission types for the various data rates.

Table 3-3 802.11g Modulation and Transmission Types

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


IEEE 802.11a OFDM Physical Layer

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

IEEE 802.11a Channels

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

Figure 3-3 Channel Set Example



For the US-based 802.11a/n/ac standard, the 5 GHz unlicensed band covers 300 MHz of spectrum and supports 12 channels. As a result, the 5 GHz band is actually a conglomerate of three bands in the United States:

- 5.150 to 5.250 GHz (UNII-1)
- 5.250 to 5.350 GHz (UNII-2)
- 5.725 to 5.875 GHz (UNII-3)

RF Power Terminology

The terms such as *dB*, *dB_i*, and *dB_m* are used to describe the amount of change in power measured at points in a system, as perceived by the radio or compared to a reference power level. The following sections cover their differences and provide general rules for their use. Effective isotropic radiated power (EIRP) is also described.

dB

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

The mathematical equation is:

$$\text{power (in dB)} = 10 * \log_{10} (\text{signal/reference})$$

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

Keep in mind that it is logarithmic, meaning that it increases or decreases exponentially and not linearly, and it is a ratio of a given value to a reference. Also keep in mind that it is multiplied by 10.

Given that it is logarithmic, there are general rules to take into consideration. An increase or decrease of 3 dB means that the signal doubled (double the power) or halved, respectively. An increase or decrease of 10 dB means that the signal went up by 10 times or down to $1/10^{\text{th}}$ of the original value.

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

dB_i

The term dB_i (dB *isotropic*) describes the forward gain of a real antenna compared with the hypothetical isotropic antenna. An isotropic antenna (a theoretical or imaginary antenna) is one that sends the same power density perfectly in all directions.

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

dB_m

The term dB_m (dB *milliwatt*) uses the same calculation as described in the dB section but has a reference value of 1 mW.

Taking into consideration the example given above in the dB section, if the power increased from 1 mW to 100 mW at the radio, the power level would increase from 0 dB_m to 20 dB_m.

Besides describing transmitter power, dB_m can also describe receiver sensitivity. Receiver sensitivity is represented as minus dB_m (-dB_m) because the signal reduces in value from its point of transmission. The sensitivity indicates the lowest power the receiver can receive before it considers the signal unintelligible.

Effective Isotropic Radiated Power (EIRP)

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

- Bridge: 20 dB_m
- 50 Foot Cable: -3.3 dB_m (negative because of cable loss)
- Dish Antenna: 21 dB_i
- EIRP: 37.7 dB_m

For more information see the Cisco TechNote *RF Power Values* at:

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


Planning for RF Deployment

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

The Cisco Prime Infrastructure provides integrated RF prediction tools that can be used to create a detailed wireless LAN design, including CAPWAP AP placement, configuration, and performance/coverage estimates. IT staff can import real floor plans into Cisco Prime Infrastructure and assign RF characteristics to various building components to increase design accuracy.

The Cisco Prime Infrastructure graphical heat maps help IT staff visualize anticipated wireless LAN behavior for easier planning and faster rollout. Cisco Prime Infrastructure also supports irregularly-shaped buildings by offering drawing tools to help organizations easily design and support WLAN deployments in such buildings. Figure 3-4 shows an example of the planning tool.

Figure 3-4 Cisco Prime Infrastructure Planning Tool



Different Deployment Types of Overlapping WLAN Coverage

How much overlapping WLAN coverage you set in the design of your wireless network depends on the usage, though with limited exceptions, all designs should be deployed to minimize retransmission and data rate shifting. Wireless networks can be deployed for data-only, voice, video, and location-based services (LBS), or a combination of all of these. The difference is in the pattern in which the APs are laid out and the amount of RF overlap in the coverage area.

When planning a WLAN deployment, consideration should be given to future uses of the WLAN deployment. Converting a WLAN deployment to support additional services beyond a data-only deployment is not simply a matter of adding APs; it can require an additional site survey and the possible relocation of existing APs.

For more information on different types of network deployments, see:


- *Cisco Unified Wireless iPhone 792x Deployment Guide:*
http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7925g/7_0/english/deployment/guide/7925dply.pdf
- *Cisco VoWLAN Troubleshooting Guide: Site Survey and RF Design Validation:*
http://www.cisco.com/en/US/docs/wireless/technology/vowlan/troubleshooting/8_Site_Survey_RF_Design_Valid.html
- *Cisco Site Survey Guide: Deploying Cisco 7920 IP Phones:*
http://www.cisco.com/en/US/docs/wireless/technology/7920/site_survey/guide/survovr.html
- *Wireless LAN Design Guide for High Density Client Environments in Higher Education:*
http://www.cisco.com/web/strategy/docs/education/cisco_wlan_design_guide.pdf
- *Cisco Wireless Mesh Access Points, Design and Deployment Guide:*
http://www.cisco.com/en/US/docs/wireless/technology/mesh/7.3/design/guide/Mesh_chapter_0100.html
- *Wi-Fi Location-Based Services Design Guide: Best Practices Location-Aware WLAN Design Considerations:*
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/wifich5.html>

Data-Only Deployments

Data-only deployments do not require a large amount of overlap. This is because 802.11 clients respond to a lower signal from a nearby AP by stepping down their rate and extending the time to transmit. The required overlap is determined by the WLAN data rate requirement described in [WLAN Data Rate Requirements, page 3-13](#). For data-only networks a general rule for the separation of APs is typically 120 to 130 feet. But when making your estimation for AP separation be sure to factor in objects that affect RF coverage such as wall densities, machinery, elevators, or even open spaces with steel cages. Your results can vary depending on the RF environment. Radio resource management (RRM, also known as *Auto-RF* in WLC) has been developed for this type of deployment and is very useful for controlling the RF coverage.

Voice Deployments

[Figure 3-5](#) shows the voice network pattern and overlap.

Figure 3-5 Single Floor Site Survey for Voice

For a voice network the APs are grouped closer together and have more overlap than a data-only installation because voice clients need to roam to a better AP before dropping packets. Generally, you should create smaller cells than for data-only networks and ensure the overlapping cell edges are at or above -67 dBm. This accomplishes a number of things including greater homogeneity across a single cell and reduced processor load in the client device, which increases link stability and reduces latency. Although only one AP might be required for a defined area, Cisco recommends that you have two APs on non-overlapping channels with a received signal strength indication (RSSI) above 35 at all times in your installation for the purposes of redundancy and load balancing. For example, for a Cisco 792x VoIP phone deployment, Cisco recommends that you have an RSSI above 35 at all times in your installation. This is to ensure that the phone has good reception as well as allow some over-subscription. It also enhances roaming choices for the phone.

Keep in mind that designing for low noise background is as important as relatively high energy density within cells. This means that a good baseline power setting for the AP is in the 35 to 50 mW range. This generally requires approximately 15 percent more APs than if you deployed a coverage model at 100 mW.

Pre-site surveys are useful for identifying and characterizing certain challenging areas and potential sources for interference such as existing WLANs, rogue clients, and non-802.11 interference from sources such as microwave ovens and certain cordless telephones. Following a design that is reviewed and approved by all stakeholders, post-site surveys can be considered as an excellent audit mechanism to ensure that the coverage model complies with the intended functional requirements as set forth by the stakeholders. For more information on site surveys, see: *Cisco VoWLAN Troubleshooting Guide: Site Survey and RF Design Validation*:

http://www.cisco.com/en/US/docs/wireless/technology/vowlan/troubleshooting/8_Site_Survey_RF_Design_Valid.html

When making your estimation for separation, remember to factor in objects that affect RF coverage such as wall densities, machinery, elevators, or even open spaces with steel cages. Your results can vary depending on the RF environment. Be sure to include transient dynamics such as forklifts, large groups of people, or large objects moved through the area by crane or similar load bearing devices.

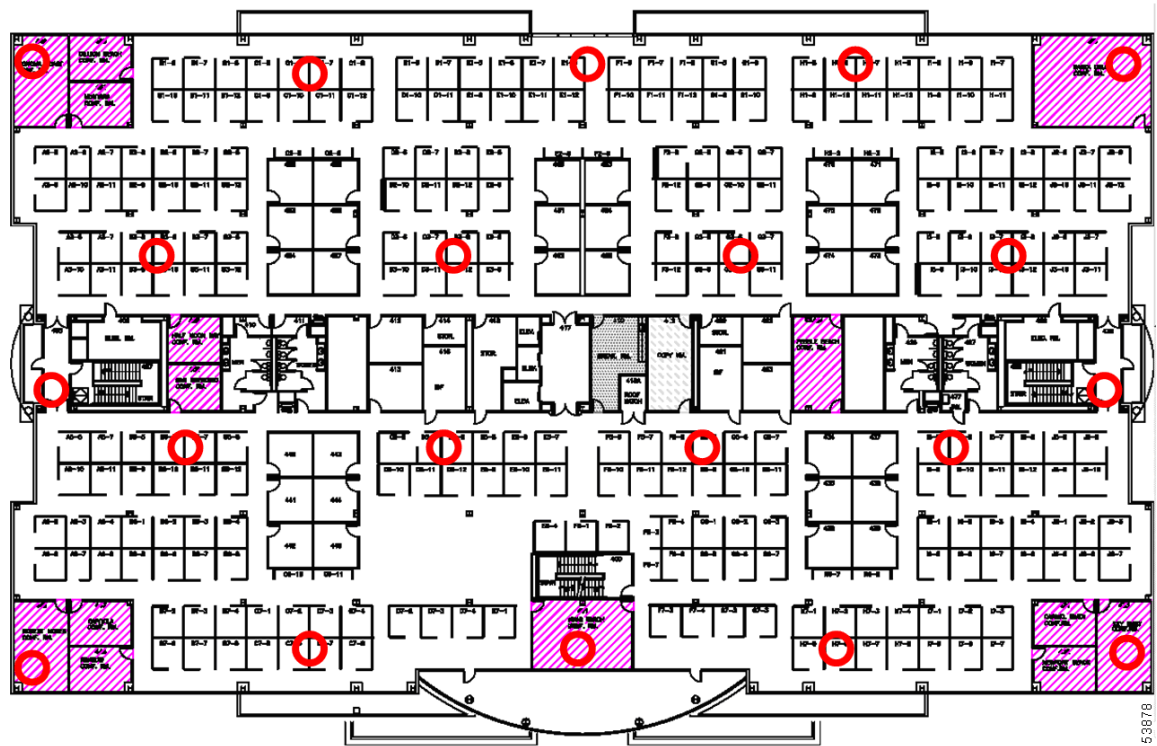
The WLC can be used to create a preliminary site evaluation (*predictive site survey*) that allows for the rapid deployment of a WLAN infrastructure and for the making of RF measurements of the area. A hand-walked site survey is also effective insurance for complex areas such as those commonly found in health care, retail, and manufacturing. For more information on a wireless voice deployments, see [Chapter 9, “VoWLAN Design Recommendations.”](#)

Location-Based Services Deployments

The last type of deployment to discuss is about location-based services (LBS), which could be the most complex of the network applications because it relies not only on excellent cell coverage but also on the optimal location of APs. LBS deployments can simultaneously track thousands of devices by using the WLAN infrastructure. Examples in LBS include Wi-Fi tag type deployments or asset tracking deployments to locate equipment or devices by way of the wireless network. LBS can also be used simply to indicate where wireless clients are throughout the wireless network in relation to a drawing or diagram. This information can be used to increase wireless infrastructure security by providing the location of rogue clients or APs. One thing it does is greatly improve client troubleshooting capabilities.

APs are laid out in a staggered pattern for location-based services deployments. The staggered pattern allows for more accurate estimation of the location of a device. [Figure 3-6](#) shows a typical pattern.

Figure 3-6 Example of a Single Floor Location Management Deployment



For a discussion of location-based services, see [Chapter 11, “Cisco Mobility Services Engine,”](#) and the Cisco *Wi-Fi Location Based Services 4.1 Design Guide* at:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/wifich1.html>.

The Cisco 7921G and the Cisco 7920 are Cisco VoWLAN handsets. Their use is one of the most common reasons for deploying QoS on a WLAN.

For more information on the 7920 and 7921G handsets, see:

- *Cisco Unified Wireless IP Phone 7921G:*
http://www.cisco.com/en/US/products/hw/phones/ps379/products_data_sheet0900aecd805e315d.html
- *Cisco Unified Wireless IP Phone 7920:*
http://www.cisco.com/en/US/products/hw/phones/ps379/products_data_sheet09186a00801739bb.html

Deploying VoWLAN infrastructure involves more than simply providing QoS on WLAN. When planning a voice WLAN deployment you need to consider site survey coverage requirements, user behavior, roaming requirements, and admission control. These are covered in the following guides:

- *Design Principles for Voice Over WLAN:*
http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/net_implementation_white_paper0900aecd804f1a46.html
- *Cisco Unified Wireless IP Phone 7921G Administration Guide:*
http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7921g/5_0_1/english/administration/guide/21adm501.html
- *Cisco Wireless IP Phone 7920 Design and Deployment Guide:*
http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7920/5_0/english/design/guide/7920ddg.html


WLAN Data Rate Requirements

Data rates affect AP coverage areas. Lower data rates (such as 1 Mbps) can extend the coverage area farther from the AP than higher data rates (such as 54 Mbps) as illustrated in [Figure 3-7](#) (not drawn to scale). Therefore, the data rate (and power level) affects coverage and consequently the number of APs required for the installation for different data rates, as illustrated in [Figure 3-8](#). As part of the planning process, consider the required data rates, the required range, and the required reliability.

Data Rates Compared to Coverage Area

Different data rates are achieved by the AP using different encoding techniques on the wireless link allowing data to be more easily recovered from noise; this can be seen in the different receiver sensitivities for the different data rates. The number of *symbols* (groups of chips) sent out for a packet at the 1 Mbps data rate is greater than the number of symbols used for the same packet at 11 Mbps. This means that sending data at the lower bit rates requires more time than sending the equivalent data at a higher bit rate. And when there is more than one client associated to the radio the lower rate client affects the maximum data throughput of higher rate clients by taking longer to transmit a packet of the same length.

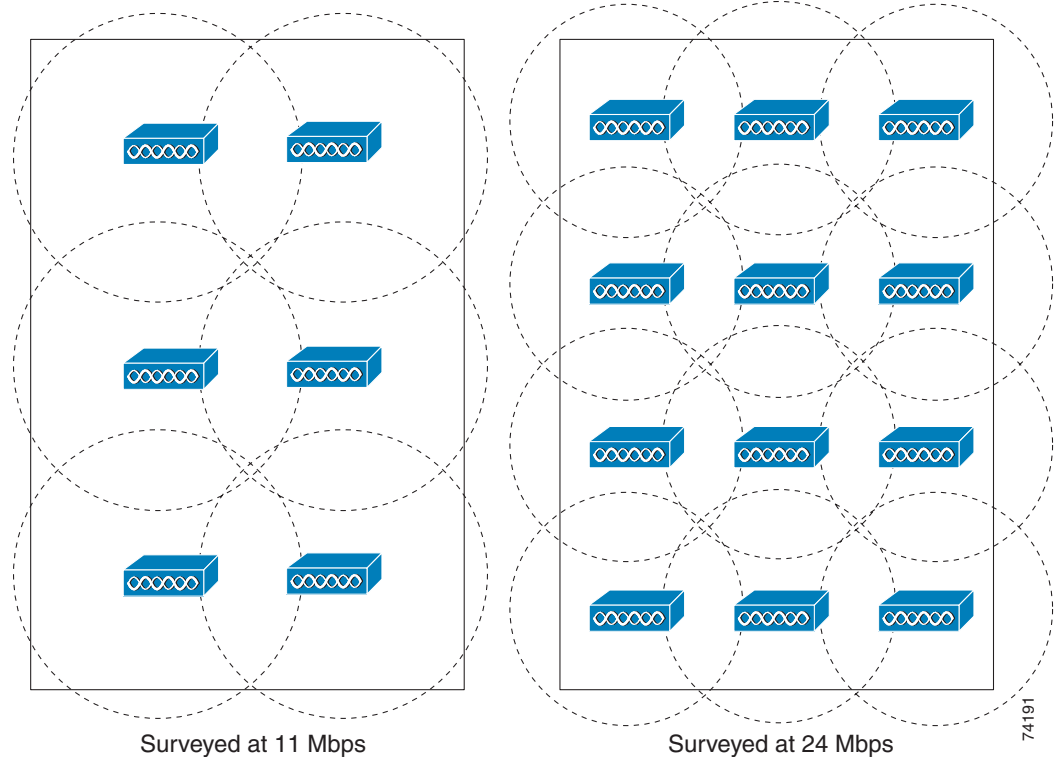
The actual diameter of the coverage, as shown in [Figure 3-7](#), depends on factors such as environment, power level, and antenna gain.

Figure 3-7 Data Rate Compared with Coverage

For example, deployed indoors using the standard antennas on the NIC card and APs, the diameter of the 1 Mbps circle is approximately 700 feet (210 meters), and the diameter of the 11 Mbps circle is about 200 feet (60 meters). These coverage diameters depend upon the type of indoor environment. An open office plan building is different from one with offices and solid walls. Increasing the gain of the antenna can increase the distance and change the shape of the radiation pattern to be focused in specific directions rather than being radiated evenly.

AP Density for Different Data Rates

The minimum required reliable data rate has a direct impact upon the number of APs needed in the design, along with power setting, antenna gain, and location. [Figure 3-8](#) shows coverage comparison and AP density for different data rates. Although six APs with a minimum data rate of 11 Mbps might adequately service an area, it might take twice as many APs to support a minimum data rate of 24 Mbps, and more again to support a minimum data rate of 48 Mbps for the same coverage area.

Figure 3-8 Coverage Comparison and AP Density for Different Data Rates

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

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

- Broadcast and multicast (if enabled) are sent at the lowest associated data rate to ensure that all clients can receive the packets. This reduces the throughput of the WLAN because traffic must wait until frames are processed at the slower rate.
- Clients that are farther away, and therefore accessing the network at a lower data rate, decrease the overall throughput by causing delays while the lower bit rates are being serviced. It might be better to force the clients to roam to a closer AP so as not to impact the performance of the rest of the network.
- If a 54 Mbps service is specified and provisioned with APs to support *all* data rates (for example), clients at lower rates can associate with the APs that can create a coverage area greater than planned, thereby increasing the security exposure (by allowing association from outside the building) and potentially interfering with other WLANs.

Client Density and Throughput Requirements

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


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

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

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

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

You can increase the potential per-user throughput by decreasing the number of users contending for the aggregate throughput provided by a single AP. This can be done by decreasing the size of the coverage area or adding a second AP on a non-overlapping channel in the same coverage area. To reduce the coverage area the AP power or antenna gain can be reduced resulting in fewer clients in that coverage area. This means you need more APs for the same overall area, thereby increasing the cost of deployment. An example of this is shown in [Figure 3-9](#).

Figure 3-9 *Changing the Output Power to Increase Client Performance***Note**

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

WLAN Coverage Requirements

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

Power Level and Antenna Choice

Power level and antenna choice go hand-in-hand to determine AP placement. Together, these two variables determine where and how powerful the RF is in any given place in the environment. Along with choosing the correct antenna to produce the required coverage area, Cisco recommends that you use RRM to control the power level and provide the optimal channel/power plan. For more information, see [Radio Resource Management, page 3-27](#).

An antenna gives the wireless system three fundamental properties:

- **Gain**—A measure of increase in power introduced by the antenna over a theoretical (isotropic) antenna that transmits the RF energy equally in all directions.
- **Direction**—The shape of the antenna transmission pattern. Different antenna types have different radiation patterns that provide various amounts of gain in different directions.
- **Polarization**—Indicates the direction of the electric field. An RF signal has both an electric and magnetic field. If the electric field is orientated vertically, the wave is said to be vertically polarized.

A good analogy for how an antenna works is the reflector in a flashlight. The reflector concentrates and intensifies the light beam in a particular direction similar to what a parabolic dish antenna does to an RF source in a radio system.

Gain and direction mandate range, speed, and reliability while polarization affects reliability and isolation of noise.

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


http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/product_data_sheet09186a008008883b.html

Omni-Directional Antennas

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

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

For example, the omni-directional antenna shown in [Figure 3-10](#) shows an omni-directional antenna RF radiation pattern in the vertical and horizontal direction. This is an actual measurement so it does not follow the donut lines perfectly, but does show from where this shape comes. As described above, other RF-affecting variables (people in the room, amount of devices stored in the facility, leaves on trees for outdoor deployment, interference from different RF sources, and so on) can affect the real RF coverage pattern.

Figure 3-10 Omni-Directional RF Pattern

With regards to the pattern in [Figure 3-10](#), this might not be the right antenna to use on a wall especially if it is mounted along an exterior wall where the pattern can radiate outside of the building.


Patch Antennas

The patch antenna is a type of directional antenna. Patch antennas not only radiate away from the wall or place where they are mounted, but also have rear and side lobes that produce a weakened but still potentially useful RF region. [Figure 3-11](#) shows the real horizontal pattern of a patch wall mount antenna. Although most of the coverage area is in front of the patch antenna, notice the back and side RF pattern from the center area. Again, antenna selection is important because it defines the radiation pattern and where wireless connectivity is possible.

Figure 3-11 Patch Wall Mount Antenna Horizontal Plane

Understanding Antenna Patterns

Patch (Directional)




346821

Dipole Antennas

Dipole antennas (see [Figure 3-12](#)) are the most basic type of radio antennas. They come in various geometries with different feeding mechanisms and radiating elements. Dipole antennas are the simplest and most practical antennas from a theoretical point of view and as of today, probably the most common antenna type.

Figure 3-12 Dipole Antenna



348820

Security Policy Requirements

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

RF Environment

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

- 2.4 GHz cordless phones and Bluetooth
- Walls fabricated from wire mesh and stucco
- Filing cabinets and metal equipment racks
- Wireless cameras
- Heavy duty electric motors, welders, robots, and things that potentially make sparks
- Fire walls and fire doors as they contain additional metal that reflects radio waves
- Concrete

- Refrigerators
- Air conditioning duct-work
- Other radio equipment such as amateur (ham) radios
- Microwave ovens, especially industrial paint drying equipment
- HVAC ducts near antennas can cause directionality or increased retries
- Large transient elements such as forklifts or metal fabrications
- Other WLAN equipment that is not part of your equipment (for example, a nearby company)

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

RF Deployment Best Practices

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

- Cisco recommends that for a given AP the number of users per AP be:
 - 15 to 25 for data-only users
 - 7 to 8 voice users (using Cisco 792x VoIP wireless handsets or similar) when data is present

This number should be used as a guideline and can vary depending on the handset in use. Check your handset requirements.

- The AP data rates should be limited to those designed and for which the site survey was performed. Enabling lower data rates can cause increases in co-channel interference and greater throughput variations for clients.
- The number of APs depends on coverage and throughput requirements, which can vary. For example, the Cisco Systems internal information systems (IS) group currently uses six APs per 3000 square feet of floor space for data-only operation.



Note

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

Manually Fine-Tuning WLAN Coverage

A number of factors can affect the WLAN coverage. They include:

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

For a given data rate and location, the WLAN designer can alter power levels and/or elect to use a different antenna to effect changes to the coverage area and/or coverage shape. Altering power levels or channel selection can be done manually as described below, or Cisco Prime Infrastructure can do this

automatically by way of the RRM algorithms. Cisco recommends that you use RRM to control the power level and channel, keeping in mind that the channel changing algorithm is highly dampened so that only a very disruptive (and persistent) interference source can cause a change to the channel topology. This can then cause clients to reassociate and voice calls to be dropped. Changes in AP power do not impact clients. See [Figure 3-16](#) for more details.

Channel and Data Rate Selection

Channel selection depends on the frequencies that are permitted for a particular region. For example, the North American and ETSI 2.4 GHz band permits allocation of three non-overlapping channels (1, 6, and 11), while the 5 GHz band permits 23 channels.


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

- Overlapping cells should use non-overlapping channels.
- Where channels must be re-used in multiple cells, those cells should have minimal overlap with each other. [Figure 3-13](#) shows this overlap pattern. In 802.11a/n/ac deployments, adjacent channels should be avoided as overlapping cells.

Recommendations for Channel Selection

[Figure 3-13](#) shows an example of a typical 2.4-GHz band channel configuration. While channel selection is typically done automatically, it can also be done manually as described in [Manual Channel Selection](#), page 3-24.

Figure 3-13 Channel Allocated To APs



A site survey should be conducted using the same frequency plan as intended for the actual deployment. They are used to test the environment for connectivity and best AP placement. Some sites have high noise backgrounds that could prohibit the use of one or more channels. The survey provides a better estimate of how a particular channel at a particular location will react to the interference and the multipath. Channel selection also helps in planning for co-channel and the adjacent channel interference, and provides information about where you can reuse a frequency (see [Figure 3-14](#)).

**Note**

For more information on site surveys, see the publication *802.11 Wireless Network Site Surveying and Installation* (Cisco, 2005), by Bruce E. Alexander.

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

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

Manual Channel Selection

Figure 3-14 is a screenshot of the Cisco WLC window for configuring one of the 802.11b radios under the wireless selection. On the top-right side, channel 11 has been manually selected and the transmit power is set to 1, the highest level (8 is the lowest).


**Note**

The assignment method should normally be left at the global setting unless there is a desire to manually control these settings. This allows the WLC to dynamically change the channel number and transmit power as determined by RRM. See [Radio Resource Management, page 3-27](#) for more information.

Figure 3-14 Channel Assignment

221934

It is also possible to implement a dual band deployment scheme, as shown in [Figure 3-15](#). The top left portion of the diagram shows the 802.11b/g/n-only deployment, which uses the three non-overlapping channels (channels 1, 6, and 11) to map out a pattern that has the least co-channel interference; that is, interference from an AP close by that is on the same channel that is operating at sufficient power levels with its coverage pattern overlapping with that of another access point. It also shows an 802.11a/n/ac deployment, which uses the eight non-overlapping channels. The right side of the diagram illustrates how the channels would be mapped in a dual band deployment.

Figure 3-15 Dual Band Deployment Diagram

346822

Data Rate Selection

[Figure 3-16](#) is a screenshot of the Cisco WLC window for managing the global 802.11b/g/n parameters. The data rate settings are shown on the right side of the screen.

Figure 3-16 Data Rate Assignment

The screenshot shows the Cisco Wireless Configuration interface for 802.11b/g Global Parameters. The left sidebar contains a tree view with options like Access Points, Mesh, Rogues, Clients, 802.11a/n, 802.11b/g/n, Network, RRM, Auto RF, DCA, Client Roaming, Voice, Video, High Throughput (802.11n), Country, and Timers. The main content area is divided into three sections: General, Data Rates, and CCX Location Measurement.

General

802.11b/g Network Status	<input checked="" type="checkbox"/> Enabled
802.11g Support	<input checked="" type="checkbox"/> Enabled
Beacon Period (milliseconds)	100
DTIM Period (beacon intervals)	1
Short Preamble	<input checked="" type="checkbox"/> Enabled
Fragmentation Threshold (bytes)	2346
Pico Cell Mode	<input type="checkbox"/> Enabled
DTPC Support	<input checked="" type="checkbox"/> Enabled

CCX Location Measurement

Mode	<input type="checkbox"/> Enabled
------	----------------------------------

Data Rates**

1 Mbps	Disabled
2 Mbps	Disabled
5.5 Mbps	Disabled
6 Mbps	Disabled
9 Mbps	Disabled
11 Mbps	Disabled
12 Mbps	Mandatory
18 Mbps	Supported
24 Mbps	Supported
36 Mbps	Supported
48 Mbps	Supported
54 Mbps	Supported

** Data Rate 'Mandatory' implies that clients who do not support that specific rate will not be able to associate. Data Rate 'Supported' implies that any associated client that also supports that same rate may communicate with the AP using that rate. But it is not required that a client be able to use the rates marked supported in order to associate.

2218935

Data Rate Modes

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

Data rates are expressed in megabits per second. You can set the mode of each data rate as mandatory, supported, or disabled.

Mandatory Mode

Mandatory mode allows transmission for all packets, both unicast and multicast. The data rate on at least one of the APs must be set to Mandatory, and all clients that associate to the AP must be able to physically support this data rate on their radio to use the network. Additionally, for the wireless clients to associate to the AP they must be able to currently receive packets at the lowest mandatory rate and their radios must physically support the highest mandatory data rate. If more than one data rate is set to mandatory, multicast and broadcast frames are sent at the highest common mandatory transmission rate of all associated clients (the lowest mandatory receive rate of all of the clients). This allows all clients to receive broadcast packets. The lowest mandatory rate is normally set at 1 Mbps.

Supported Mode

Supported mode allows transmission for unicast packets only. The AP transmits only unicast packets at this rate; multicast and broadcast packets are transmitted at one of the data rates set to Mandatory. The wireless clients always attempt to transmit and receive at the highest possible data rate. They negotiate

with the AP for the highest data rate set to supported or mandatory to transmit and receive unicast packets. The wireless client devices are able to receive broadcast or multicast packets at any mandatory rate at or below the negotiated rate.

Disabled

The AP does not transmit data in this setting.

Lowest and Highest Mandatory Rate Settings

Multiple clients associated to the AP can have different transmission rates, depending on interference, obstacles, or their distance from the AP. For example, if an 802.11b client is far from the AP and can only transmit and receive at a speed of 1 Mbps because of the distance, it would be able to associate to the AP because the lowest mandatory rate (see [Figure 3-16](#)) is set to 1 Mbps. If a second 802.11g client associates to the AP at 54 Mbps, the AP transmits broadcasts and multicasts at 1 Mbps because this is the highest mandatory rate that all clients can receive. If the lowest mandatory rate was set to 5.5 Mbps, the 802.11b client would not be able to associate to the AP because it could not receive broadcast packets at the lowest mandatory rate.

In [Figure 3-16](#), note that the highest mandatory setting is 11 Mbps. The highest mandatory rate tells the AP at what rate the client radio must be able to transmit physically. This does not mean that they are actually transmitting and receiving packets at that rate; it just means that the radio physically supports that rate. The wireless client needs only to be able to receive packets at the lowest mandatory rate. 802.11b devices are able to associate to the AP shown in [Figure 3-16](#) because their radios can physically transmit at 11 Mbps. If a higher data rate (such as 18 Mbps) is set to mandatory, only 802.11g clients are able to associate to the APs.

Setting any of the OFDM rates (rates above 1 Mbps) to mandatory disables 802.11b connectivity. This can, for example, allow the administrator to exclude 802.11b clients from the AP by requiring an 802.11g data rate or setting a minimum transmission rate of all clients by disabling 802.11 rates. The reason this might be done is that the same 1500 byte packet at a lower data rate takes a longer time to transmit and, thus, lowers the effective data rate for all wireless clients associated to the AP.

Radio Resource Management

In the Cisco WLAN *split MAC* architecture (see [Chapter 2, “Cisco Unified Wireless Technology and Architecture”](#)), the processing of 802.11 data, management protocols, and AP capabilities is distributed between a CAPWAP AP and a centralized WLC. More specifically, time-sensitive activities such as probe response and MAC layer encryption are handled by the APs while all other functions are sent to the WLCs where system-wide visibility is required.

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



Note

An RF network group (or RF group) is not the same as a mobility group. A mobility group defines a mobility domain of 1 to 24 WLCs in which a client would not be required to change IP address during a roaming event. This is accomplished by creating Ethernet over IP tunnels for forwarding client data from an *anchor* controller to the *foreign* controller handling the new AP servicing the client.

Radio resource management (RRM) algorithms adjust the channel and power (using dynamic channel assignment and dynamic transmit power control, respectively) to maintain the RF coverage area. RRM adjusts the power level of the AP to maintain a baseline signal strength with neighboring APs at (configurable) -65 dBm (see [Overview of RRM Operation](#)). RRM adjusts the channel of the AP when it notices nearby interference sources on the channel that the AP is currently located. RRM continues to optimize the RF coverage for the best reception and throughput for the wireless network.

**Note**

The transmit power control and dynamic frequency management performed by RRM are not the TPC and DFS required for operation in the UNII-2 bands that are defined in 802.11h.

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

Overview of RRM Operation

If group mode is enabled in the WLCs, they select a leader in each RF group and form an RF domain. The function of the leader is to collect the network-wide neighbor data packets for the APs from the group of WLCs and do channel/power computation for an optimal system-wide map. If group mode is not enabled, the controllers run computations based only on the neighbor data gathered from the APs connected by way of CAPWAP, trying to optimize the signal to -70 dBm between APs.

The APs transmit RRM neighbor data packets at full power at regular intervals (neighbor data packets include information about the signal strength and radio elements in the environment). These messages contain a field that is a hash of the RF network (group) name, BSSID, and time stamp. The APs accept only RRM neighbor packets sent with the default RF network name (*RF Network Name*).

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

[Table 3-4](#) provides a summary of the various functions of the devices in the system.

**Note**

TPC performs only downward power level adjustments. Coverage hole detection and correction increases power levels on APs.

Table 3-4 **Device Function**

Device	Functions
RF Group Leader	WLC that collects AP neighbor data from the other WLCs in the RF group and analyzes it for system-wide TPC and DCA. TPC adjusts power levels downward.

Table 3-4 **Device Function (continued)**

Local WLC	Collects data and runs the Coverage Hole Detection and Correction algorithm. Adjusts power levels upward if necessary for clients
CAPWAP AP	<ul style="list-style-type: none"> • Sends neighbor messages on all channels at full power at configured interval • Verifies neighbor hash on received neighbor messages • Scans configured channels for noise, interference, IDS/rogue detection, and alerts

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

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

RRM Configuration Settings

Radio resource management (RRM) also called as *Auto-RF* in the WLC, can be turned on or off in the WLC by way of a global setting of Channel Selection (see [Figure 3-14](#)). From this window you can also manually set the channel and transmit level for the APs. Additionally, RRM can be turned on or off from the WLC global Auto-RF configuration window. Keep in mind that RRM is performed on a per band basis. RF group computations for the 5 GHz band are separate from the computations for the 2.4 GHz band.

The Auto-RF configuration window is divided into three sections that you move between using the scroll bar on the right. The first section (see [Figure 3-17](#)) is for dynamic channel assignment. This allows the WLC to automatically change the channel that the AP is on (for more information, see [Dynamic Channel Assignment](#)).

Figure 3-17 Auto-RF (Section 1)

The screenshot shows the Cisco Wireless LAN Controller configuration page for 802.11a Global Parameters > Auto RF. The page is divided into two main sections: RF Group and RF Channel Assignment. The RF Group section includes settings for Group Mode (Enabled), Group Update Interval (600 secs), Group Leader (00:0b:85:40:98:40), Is this Controller a Group Leader? (Yes), and Last Group Update (46 secs ago). The RF Channel Assignment section includes settings for Channel Assignment Method (Automatic, Every 600 sec), Avoid Foreign AP interference (Enabled), Avoid Cisco AP load (Enabled), Avoid non-802.11a noise (Enabled), Signal Strength Contribution (Enabled), Channel Assignment Leader (00:0b:85:40:98:40), and Last Auto Channel Assignment (46 secs ago). A red box highlights the RF Channel Assignment section.

**Note**

For more information, see the *Cisco Wireless LAN Controller Configuration Guide*, http://www.cisco.com/en/US/docs/wireless/controller/7.3/configuration/guide/b_cg73.html

The first group of settings in the Auto-RF configuration window, the RF Group, is used to determine whether the WLC joins the dynamic grouping with the other WLCs in the group. Dynamic grouping helps the WLC find out about APs that are neighbors but might be associated to another WLC in the mobility group. If this is disabled, the WLC only optimizes the parameters of the access points that it knows about (that is, the ones that are associated to it). The group leader indicates the MAC address of the elected leader. You can find the MAC address of the controller in the WLC Inventory window by clicking on the *Controller* menu at the top and then selecting *Inventory*.

The second section is for assigning the transmit (Tx) power level (Figure 3-18). The power level can be fixed for all APs, or it can be automatically adjusted. This window also indicates the number of neighbors the AP has and the power thresholds for which it is adjusting.

Figure 3-18 Auto-RF (Section 2)

CISCO MONITOR WLANs CONTROLLER **WIRELESS** SECURITY MANAGEMENT COMMANDS HELP

Save Configuration | Ping | Logout | Refresh

Wireless

- Access Points
 - All APs
 - Radios
 - 802.11a/n
 - 802.11b/g/n
 - AP Configuration
- Mesh
- Rogues
- Clients
- 802.11a/n
 - Network
 - RRM
 - Auto RF
 - DCA
 - Pico Cell
 - Client Roaming
 - Voice
 - Video
 - DFS (802.11h)
 - High Throughput (802.11n)
- 802.11b/g/n
- Country
- Timers

Tx Power Level Assignment

Power Level Assignment Method: ☒ Automatic Every 600 sec
☐ On Demand [Invoke Power Update now](#)
☐ Fixed

Power Threshold: -65 dBm
 Power Neighbor Count: 3
 Power Update Contribution: SNI.
 Power Assignment Leader: 00:0b:85:40:98:40
 Last Power Level Assignment: 46 secs ago

Profile Thresholds

Interference (0 to 100%)	<input type="text" value="10"/>
Clients (1 to 75)	<input type="text" value="12"/>
Noise (-127 to 0 dBm)	<input type="text" value="-70"/>
Coverage 3 to 50 dBm)	<input type="text" value="16"/>
Utilization (0 to 100%)	<input type="text" value="80"/>
Coverage Exception Level (0 to 100 %)	<input type="text" value="25"/>
Data Rate 1 to 1000 Kbps	<input type="text" value="1000"/>
Client Min Exception Level (1 to 75)	<input type="text" value="3"/>

The third section (Figure 3-19) is for profile thresholds.

Figure 3-19 Auto-RF (Section 3)

CISCO MONITOR WLANs CONTROLLER **WIRELESS** SECURITY MANAGEMENT COMMANDS HELP

Save Configuration | Ping | Logout | Refresh

Wireless

- Access Points
 - All APs
 - Radios
 - 802.11a/n
 - 802.11b/g/n
 - AP Configuration
- Mesh
- Rogues
- Clients
- 802.11a/n
 - Network
 - RRM
 - Auto RF
 - DCA
 - Pico Cell
 - Client Roaming
 - Voice
 - Video
 - DFS (802.11h)
 - High Throughput (802.11n)
- 802.11b/g/n
- Country
- Timers

Profile Thresholds

Interference (0 to 100%)	<input type="text" value="10"/>
Clients (1 to 75)	<input type="text" value="12"/>
Noise (-127 to 0 dBm)	<input type="text" value="-70"/>
Coverage 3 to 50 dBm)	<input type="text" value="16"/>
Utilization (0 to 100%)	<input type="text" value="80"/>
Coverage Exception Level (0 to 100 %)	<input type="text" value="25"/>
Data Rate 1 to 1000 Kbps	<input type="text" value="1000"/>
Client Min Exception Level (1 to 75)	<input type="text" value="3"/>

Noise/Interference/Rogue Monitoring Channels

Channel List:

Monitor Intervals (60 to 3600 secs)

Noise Measurement	<input type="text" value="180"/>
Load Measurement	<input type="text" value="60"/>
Signal Measurement	<input type="text" value="60"/>
Coverage Measurement	<input type="text" value="180"/>

Factory Default

Set all Auto RF 802.11a parameters to Factory Default.
[Set to Factory Default](#)

Sample 'show ap auto-rf' Command Output

The WLC analyzes the information passed to it by the APs and determines a pass/fail status for each of these thresholds. These pass/fail profiles are best seen in the output of the following WLC CLI **show ap auto-rf radio ap_name** command. This command displays RF statistics from the radio being used.



Note

The same information can be seen in graphical form in the WLC Monitor -> **802.11b/g/n Radios -> Detail** window.

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

```

Last Channel Change Time . . . . . Thu Mar 9 12:18:03 2006
Recommend Best Channel . . . . . 11
RF Parameter Recommendations
Power Level . . . . . 1
RTS/CTS Threshold . . . . . 2347
Fragmentation Threshold . . . . . 2346
Antenna Pattern . . . . . 0

```

The following sections describe some of the WLC RRM settings.

Dynamic Channel Assignment

The 802.11 MAC layer uses Carrier-Sense Multiple Access/Collision Avoidance (CSMA/CA). Because of CSMA/CA-managed channel sharing, two APs on the same channel in the same vicinity get approximately half the capacity of what two APs would get on different channels in the same vicinity. This is caused by the 802.11 MAC sensing that the channel is busy, and deferring sending frames until the channel has become free. If the 802.11 MAC defers traffic that is not part of its own AP cell, this is considered interference. Interference from another AP on the same channel is commonly called *co-channel interference*, and is to be expected in most 2.4 GHz band 802.11 deployments. This is because there are insufficient non-overlapping channels to prevent some channel overlap from occurring. One of the goals of design, planning, and dynamic radio management is to minimize the amount of co-channel overlap, which minimizes co-channel interference and therefore maximizes AP traffic capacity. The Cisco Unified Wireless Network addresses this problem and other co-channel interference issues by dynamically allocating AP channel assignments to avoid conflict. Because the WLC, or a designated WLC (RF Group Leader), has system-wide visibility it can control how channels are *reused* to minimize co-channel interference.

The WLC examines a variety of real-time RF characteristics to efficiently handle channel assignments, including the following:

- **Noise**—Noise limits signal quality at the client and AP and can vary in range and periodicity. There are numerous potential types and effects of interference. For one, an increase in noise reduces the effective cell size. At regular intervals the WLC reassesses the RF environment of an AP and then optimizes channel selection to avoid noise sources while still maintaining overall system capacity. Channels that become unusable because of excessive noise can be avoided. If other wireless networks are present, the WLC shifts its channel usage to complement the other networks. For example, if one network is on Channel 6, the adjacent WLAN is assigned Channel 1 or 11. This increases the capacity of the network by limiting the sharing of frequencies. If a channel is used so much that no capacity is available, the WLC might choose to avoid this channel.
- **Client load**—Client load is taken into account when changing the channel structure to minimize the impact on the clients currently on the WLAN. The WLC periodically monitors the channel assignment in search of the best assignments. Change occurs only if it significantly improves the performance of the network or corrects the performance of a poorly performing AP.

The WLC combines the RF characteristic information to make system-wide decisions. The end result is an optimal channel configuration in a three-dimensional space, where APs on the floor above and below are factored into an overall WLAN configuration.

Interference Detection and Avoidance

Interference (as it pertains to a Cisco Unified Wireless Network deployment) is defined as unwanted RF signals in the same frequency band that can lead to a degradation or loss of service. These signals can either be from 802.11 or non-802.11 sources such as certain microwave ovens or certain cordless phones. It can, in certain instances, also include various sources of electromagnetic interference (EMI) such as

arc welders or radar facilities. APs constantly scan all channels for major sources of interference and regularly reports this and additional information to the WLC by way of a management link interface (CAPWAP tunnel).

If the amount of 802.11 interference hits a predefined threshold, the WLC attempts to rearrange channel assignments to optimize system performance in the presence of the interference. This might result in adjacent APs being on the same channel. Logically this represents a better scenario than staying on a channel that is otherwise totally unusable because of interference. For example, the WLC can respond to a rogue AP on channel 11 by shifting nearby APs to channel 1 or channel 6.

Dynamic Transmit Power Control

The appropriate AP transmit power settings are essential to maintaining a coverage area, not only to ensure correct (not maximum) amount of power covering an area, but also to ensure that excessive power is not used, which would add unnecessary interference to the radiating area. AP power settings are also used to control network redundancy by helping to ensure real-time failover in the event of the loss of an AP. The WLC is used to dynamically control the AP transmit power-level based on real-time WLAN conditions. In normal instances, power can be minimized to gain extra capacity and reduce interference among the APs. RRM attempts to balance APs so that they see their neighbors at -65 dBm. If an AP outage is detected, power can be automatically increased on surrounding APs to fill the coverage gap created by the loss.

RRM algorithms are designed to create the optimal user experience. For example, if the power of an AP is turned down to Level 4 (where Level 1 = highest and Level 8 = lowest) and the RSSI value of a user drops below an acceptable threshold, the AP power is increased to provide a better experience to that client. When dynamic transmit power control (DTPC) is enabled, the AP adds channel and transmit power information to *beacons* (information elements that contain information such as channel, RF power, network name, and so forth). Client devices using DTPC receive the information and adjust their settings automatically.

Coverage Hole Detection and Correction

The coverage hole detection and correction algorithm is aimed at determining coverage holes based on the quality of client signal levels and then increasing the transmit power of the APs to which those clients are associated.

The algorithm determines whether a coverage hole exists when client signal-to-noise ratio (SNR) levels pass below a given SNR threshold. The SNR threshold is considered on an individual AP basis and based primarily on the transmit power of each AP.

When the average SNR of a single client dips below the SNR threshold for at least 60 seconds, this is seen as an indication that the WLAN client does not have a viable location to roam to. The AP transmit power of that client is increased, correcting the coverage hole.

Client and Network Load Balancing

The IEEE 802.11 standard does not define the process or reasons for client roaming; therefore, it cannot be easily predicted what clients will do in any given situation. For example, all users in a conference room can associate with a single AP because of its close proximity, ignoring other APs that are farther away but with greater free capacity.

The WLC has a centralized view of client distribution across all APs. This can be used to influence where new clients attach to the network if there are multiple *good* APs available. If configured, the WLC can proactively use AP probe responses to guide clients to the most appropriate APs to improve WLAN performance. This results in a smooth distribution of capacity across an entire wireless network. Keep in mind that this load balancing is done at the time of client association, not after a client is connected.



Cisco Unified Wireless Network Architecture —Base Security Features

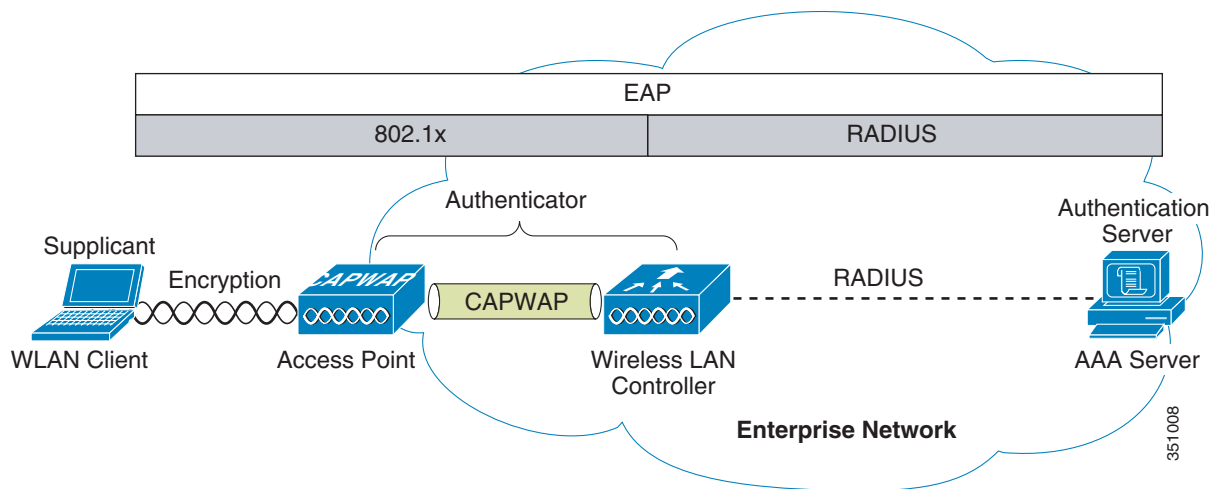
The Cisco Unified Wireless Network solution provides end-to-end security using architecture and product security features that protect wireless local area network (WLAN) endpoints, the WLAN infrastructure, and client communications.

The Cisco Unified Wireless Network solution builds upon the base security features of the IEEE 802.11-2012 standard by enhancing radio frequency (RF) and network-based security features to ensure overall security.

Secure Wireless Topology

[Figure 4-1](#) illustrates a secure wireless topology. The topology is made up of the following components with their basic roles in the 802.1X authentication process.

- WLAN client with 802.1X supplicant (wireless software) on the client
- Access point (AP) and Wireless LAN Controller (WLC) using the control and provisioning of wireless access points (CAPWAP) protocol
- RADIUS protocol carrying extensible authentication protocol (EAP) packets between the client and the authentication server
- Authentication, Authorization, and Accounting (AAA) server as the Authentication Server

Figure 4-1 Secure Wireless Topology

WLAN Security Mechanisms

Security is implemented using authentication and encryption in the WLAN network. The security mechanisms for WLAN networks are:

- Open Authentication (*no* encryption)
- Wired Equivalent Privacy (WEP)
- Cisco WEP Extensions (Cisco Key Integrity Protocol + Cisco Message Integrity Check)
- Wi-Fi Protected Access (WPA)
- Wi-Fi Protected Access 2 (WPA2)
- Cisco Adaptive Wireless Intrusion Prevention System (wIPS) with Enhanced Local Mode (ELM)

Cisco Wired Equivalent Privacy (WEP) Extensions

The original 802.11 security mechanism, WEP, is a static encryption method that applies *some* level of security and is generally viewed as insufficient for securing business communications. Cisco WLAN products addressed these insufficiencies by adopting the following enhancements to WEP:

- Cisco Key Integrity Protocol (CKIP)
- Cisco Message Integrity Check (CMIC)

These Cisco enhancements to WEP are collectively known as the Cisco WEP Extensions.

Wi-Fi Protected Access (WPA)

The 802.11 WEP standard failed to address the issue of how to manage encryption keys. The encryption mechanism itself was found to be flawed, in that a WEP key could be derived simply by monitoring client traffic. The IEEE 802.11i standard addresses these security issues found in the original 802.11 WEP standard.

WPA and WPA2 are 802.11i-based security solutions as defined by the Wi-Fi Alliance. The Wi-Fi Alliance certifies inter-operability of IEEE 802.11 products and promotes wireless LAN standards across all market segments. The Wi-Fi Alliance's test suite defines how products are tested to obtain interoperability certification with other Wi-Fi Certified products.

WPA uses Temporal Key Integrity Protocol (TKIP) for encryption and dynamic encryption key generation with either a pre-shared key or a RADIUS/802.1x-based authentication. The mechanisms introduced in WPA are designed to provide more robust security to WEP solutions without requiring a hardware upgrade.

Wi-Fi Protected Access 2 (WPA2)

WPA2 is the next generation of Wi-Fi security based on the ratified IEEE 802.11i standard and is also approved by the Wi-Fi Alliance interoperability implementation of the 802.11i standard. WPA2 provides certification in both Enterprise and Personal classifications.

The Enterprise classification requires support for a RADIUS/802.1x-based authentication and pre-shared key; Personal classification requires only a common key shared by the client and the AP.

The newer Advanced Encryption Standard (AES) mechanism introduced in WPA2 generally requires a hardware upgrade of WLAN clients and APs; however, all Cisco CAPWAP hardware is WPA2 enabled.

802.1X

802.1X is an IEEE framework for port-based access control as adopted by the 802.11i security workgroup. The framework provides authenticated access to WLAN networks.

- The 802.11 association process creates a “virtual” port for each WLAN client at the AP.
- The AP blocks all data frames apart from 802.1X-based traffic.
- The 802.1X frames carry the EAP authentication packets, which are passed through to the AAA server by the AP.
- If the EAP authentication is successful, the AAA server sends an EAP success message to the AP, where the AP then allows data traffic from the WLAN client to pass through the virtual port.
- Before opening the virtual port, data link encryption is established between the WLAN client and the AP. This is to ensure no other WLAN client can access the port established for authenticating clients.

Authentication and Encryption

The Cisco Wireless Security suite provides options to security approaches based on required or pre-existing authentication, privacy, and client infrastructure. The Cisco Wireless Security suite supports WPA, WPA2, WEP Extension, and WIPS with the ELM feature.

The following options are available:

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

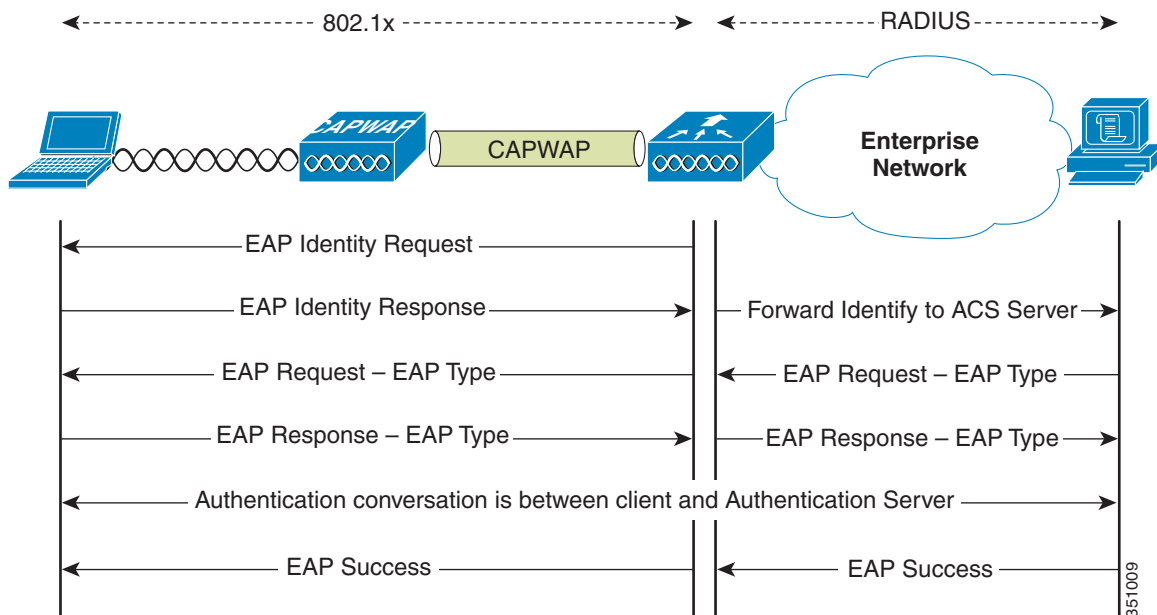
- EAP-Subscriber Identity Module (EAP-SIM)
- Encryption:
 - AES-CCMP encryption (WPA2)
 - TKIP encryption enhancements: key hashing (per-packet keying), message integrity check (MIC) and broadcast key rotation via WPA/WPA2 or WEP TKIP Cisco Key Integrity Protocol (CKIP) and Cisco Message Integrity Check (CMIC)

Extensible Authentication Protocol

Extensible Authentication Protocol (EAP) is an IETF RFC that stipulates that an authentication protocol must be de-coupled from the transport protocol. This allows the EAP protocol to be carried by transport protocols such as 802.1X, UDP, or RADIUS without making changes to the authentication protocol itself. The basic EAP protocol contains the following four packet types.

- EAP request—The request packet is sent by the authenticator to the supplicant. Each request has a type field that indicates what is being requested; for example, supplicant identity and EAP type to be used. A sequence number allows the authenticator and the peer to match an EAP response to each EAP request.
- EAP response—The response packet is sent by the supplicant to the authenticator, and uses a sequence number to match the initiating EAP request. The type of the EAP response generally matches the EAP request, except if the response is a negative-acknowledgment (NAK).
- EAP success—The success packet is sent, from the authenticator to the supplicant, when successful authentication occurs.
- EAP failure—The failure packet is sent, from the authenticator to the supplicant, when unsuccessful authentication occurs.

When using EAP in an 802.11i compliant system, the AP operates in EAP pass-through mode. Pass-through mode checks the code identifier and the length fields, and then forwards EAP packets received from the client supplicant to the AAA. EAP packets received by the authenticator from the AAA server are forwarded to the supplicant. [Figure 4-2](#) is an example of an EAP protocol flow.

Figure 4-2 EAP Protocol Flow

Authentication

Depending on your requirements, various authentication protocols such as PEAP, EAP-TLS, and EAP-FAST are used in secure wireless deployments. Regardless of the protocol, they all use 802.1X, EAP, and RADIUS as their underlying transport.

These protocols allow network access control based on the successful authentication of the WLAN client and vice-versa. This solution also provides authorization through policies communicated through the RADIUS protocol, as well as RADIUS accounting.

EAP types used for performing authentication are described in more detail below. The primary factor affecting the choice of EAP protocol is the authentication system (AAA) currently used. Ideally, a secure WLAN deployment should not require the introduction of a new authentication system, but rather should leverage the authentication systems that are already in place.

Supplicants

The various EAP supplicants that are available in the marketplace reflect the diversity of authentication solutions available and customer preferences.

Table 4-1 shows a summary of common EAP supplicants:

- **EAP-FAST**—EAP-Flexible Authentication via Secured Tunnel. Uses a tunnel similar to that used in PEAP, but does not require the use of Public Key Infrastructure (PKI).
- **PEAP MSCHAPv2**—Protected EAP MSCHAPv2. Uses a Transport Layer Security (TLS) tunnel, (the IETF standard of SSL) to protect an encapsulated MSCHAPv2 exchange between the WLAN client and the authentication server.
- **PEAP GTC**—Protected EAP Generic Token Card (GTC). Uses a TLS tunnel to protect a generic token card exchange; for example, a one-time password or LDAP authentication.

- EAP-TLS—EAP Transport Layer Security. Uses PKI to authenticate both the WLAN network and the WLAN client, requiring both a client certificate and an authentication server certificate.

Table 4-1 Comparison of Common Supplicants

	Cisco EAP-FAST	PEAP MS-CHAPv2	PEAP EAP-GTC	EAP-TLS
Single sign-on (MSFT AD only)	Yes	Yes	Yes ¹	Yes
Login scripts (MSFT AD only)	Yes	Yes	Some	Yes ²
Password change (MSFT AD)	Yes	Yes	Yes	N/A
Microsoft AD database support	Yes	Yes	Yes	Yes
ACS local database support	Yes	Yes	Yes	Yes
LDAP database support	Yes ³	No	Yes	Yes
OTP authentication support	Yes ⁴	No	Yes	No
RADIUS server certificate required?	No	Yes	Yes	Yes
Client certificate required?	No	No	No	Yes
Anonymity	Yes	Yes ⁵	Yes ⁶	No

1. Supplicant dependent

2. Machine account and machine authentication is required to support the scripts.

3. Automatic provisioning is not supported on with LDAP databases.

4. Supplicant dependent

5. Supplicant dependent

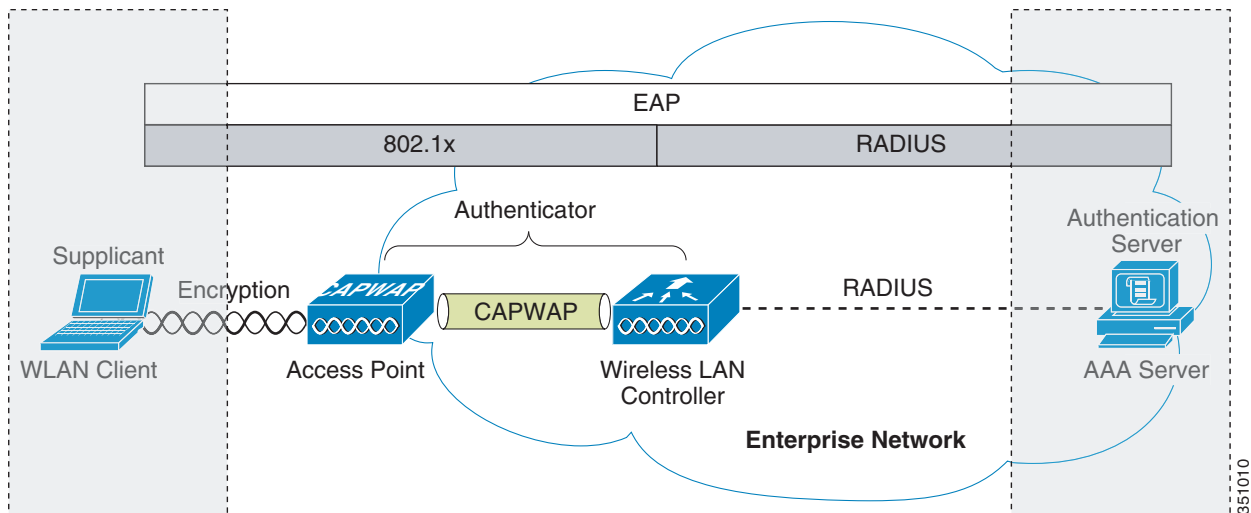
6. Supplicant dependent

Authenticator

The WLC is the authenticator acting as a relay for EAP messages exchanged between the 802.1X-based supplicant and the RADIUS authentication server. Once authentication is completed successfully, the WLC receives the following:

- A RADIUS packet containing the EAP success message
- An encryption key, which is generated at the authentication server during the EAP authentication
- RADIUS vendor-specific attributes (VSAs) for communicating policy

Figure 4-3 displays the logical location of the *authenticator* within the overall authentication architecture. The authenticator controls network access using the 802.1X protocol, and relays EAP messages between the supplicant and the authentication server.

Figure 4-3 Authenticator Location

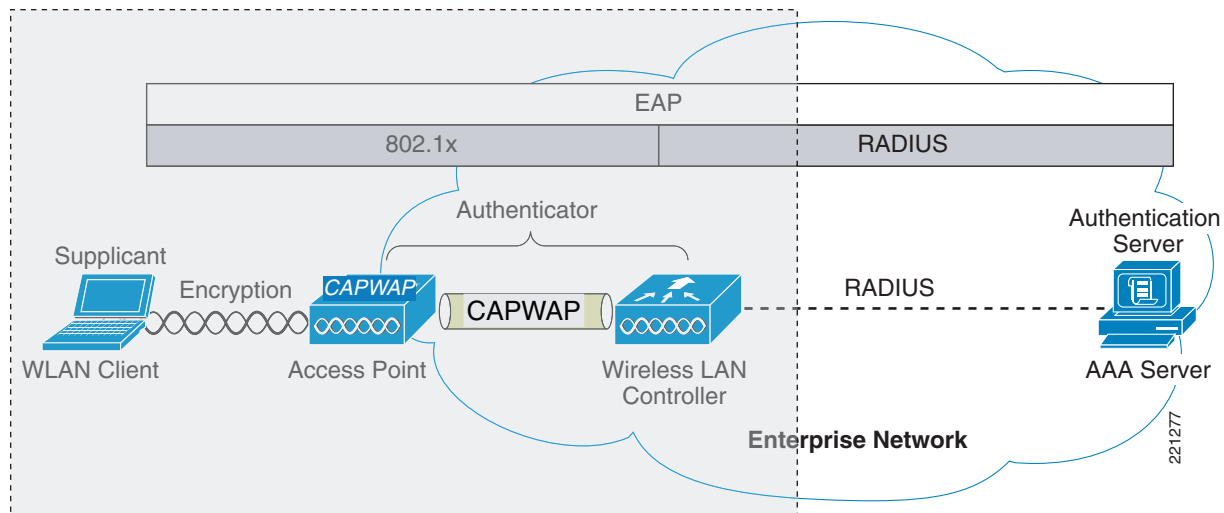
The EAP exchange sequence is as follows:

- Packet #1 is sent by the AP to the client, requesting the client identity; this begins the EAP exchange.
- Packet #2 contains the client identity, which is forwarded to the RADIUS server. Based on the client identity, in packet 2, the RADIUS server will determine to continue the EAP authentication or not.
- Packet #3 contains a RADIUS server request to use PEAP as the EAP method for authentication. The actual request depends on the EAP types configured on the RADIUS server. If the client rejects the PEAP request, the RADIUS server can offer other EAP types.
- Packets #4 through 8 are the TLS tunnel setup for PEAP.
- Packets #9 through 16 are the authentication exchange within PEAP.
- Packet #17 is the EAP message informing the supplicant and the authenticator that the authentication was successful. In addition, Packet #17 carries encryption keys and authorization information, in the form of RADIUS VSAs, to the authenticator.

Authentication Server

The authentication server used in the Cisco Secure Unified Wireless Network solution is the Cisco Access Control Server (ACS) and the Cisco Identity Services Engine (ISE). ACS and ISE are available as software that is installed on a Windows 2000 or later servers, or as an appliance. Alternatively, the authentication server role can be implemented within specific WLAN infrastructure devices such as local authentication services on an IOS AP, local EAP authentication support within the WLC, AAA services integrated in any AAA server that supports the required EAP types.

Figure 4-4 shows the logical location of the authentication server within the overall wireless authentication architecture, where it performs the EAP authentication via a RADIUS tunnel.

Figure 4-4 Authentication Server Location

After the completion of a successful EAP authentication, the authentication server sends an EAP success message to the authenticator. This message tells the authenticator that the EAP authentication process was successful, and passes the pair-wise master key (PMK) to the authenticator that is in turn used as the basis for creating the encrypted stream between the WLAN client and the AP.

Encryption

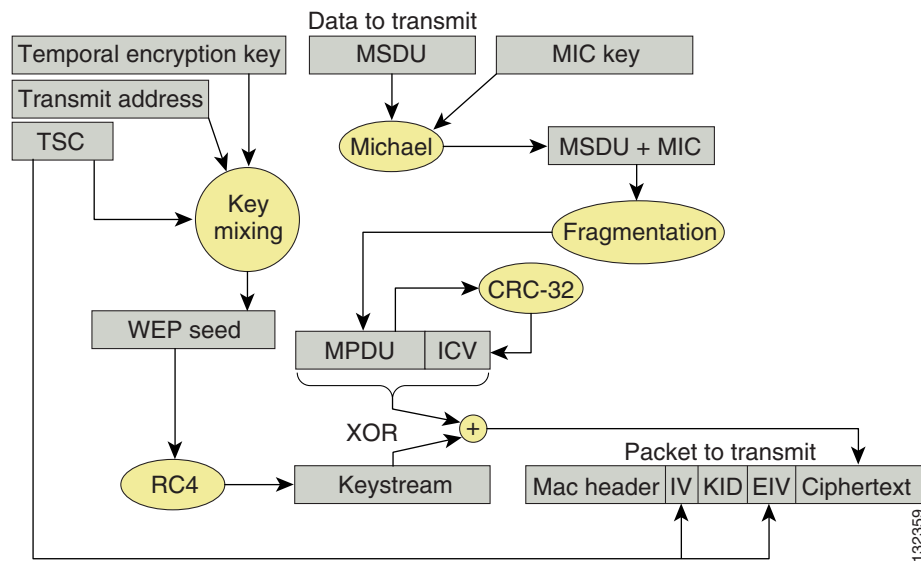
Encryption is a necessary component of WLAN security to provide privacy over a local RF broadcast network. Any new deployment should be using either TKIP (WPA/WPA2) or AES encryption.

In WPA and WPA2, the encryption keys are derived during the four-way handshake discussed later in this section.

TKIP Encryption

Enterprise-level encryption mechanisms specified by 802.11i are certified as WPA/WPA2 and WIPS by the Wi-Fi Alliance: Temporal Key Integrity Protocol (TKIP), and Advanced Encryption Standard (AES). TKIP is the certified encryption method. It provides support for legacy WLAN equipment by addressing the original flaws associated with the 802.11 WEP encryption method. It does this by making use of the original RC4 core encryption algorithm.

The hardware refresh cycle of WLAN client devices is such that TKIP is likely to be a common encryption option for a number of years to come. The AES encryption is the preferred method because it brings the WLAN encryption standards into alignment with broader IT industry standards and best practices. [Figure 4-5](#) displays a basic TKIP flow chart.

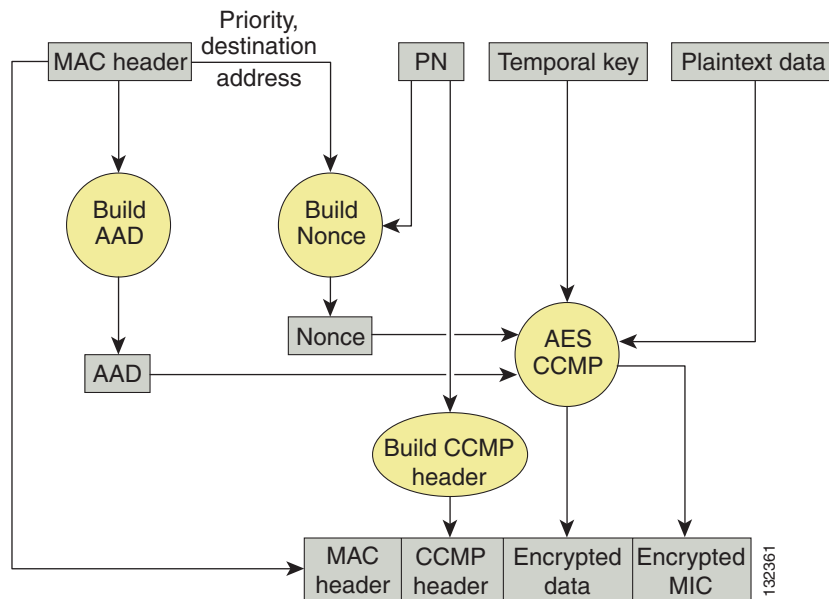
Figure 4-5 TKIP Flow Chart

The two primary functions of TKIP are the generation of a per-packet key using RC4 encryption of the MAC service data unit (MSDU) and a message integrity check (MIC) in the encrypted packet. The per-packet key is a hash of the transmission address, the frame initialization vector (IV), and the encryption key. The IV changes with each frame transmission, so the key used for RC4 encryption is unique for each frame.

The MIC is generated using the Michael algorithm to combine a MIC key with user data. The use of the Michael algorithm is a trade-off because its low computational overhead is good for performance, but it can be susceptible to an active attack. To address this, WPA includes countermeasures to safeguard against these attacks that involve temporarily disconnecting the WLAN client and not allowing a new key negotiation for 60 seconds. Unfortunately, this behavior can itself become a type of DoS attack. Many WLAN implementations provide an option to disable this countermeasure feature.

AES Encryption

Figure 4-6 displays the basic AES counter mode/CBC MAC Protocol (CCMP) flow chart. CCMP is one of the AES encryption modes, where the counter mode provides confidentiality and CBC MAC provides message integrity.

Figure 4-6 WPA2 AES CCMP

In the CCMP procedure, additional authentication data (AAD) is taken from the MAC header and included in the CCM encryption process. This protects the frame against alteration of the non-encrypted portions of the frame.

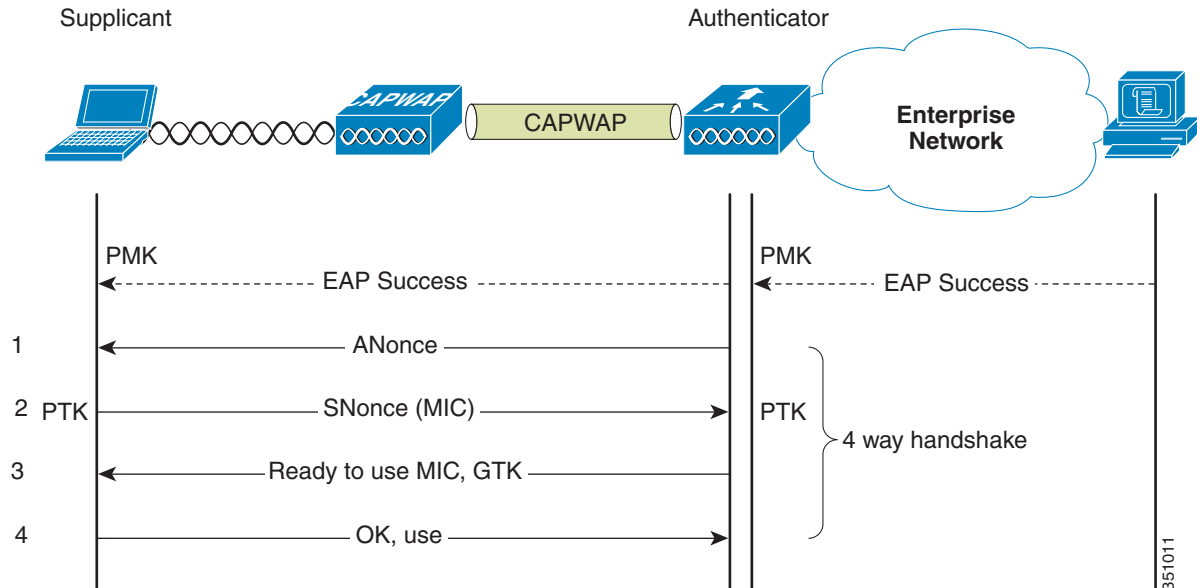
To protect against replay attacks, a sequenced packet number (PN) is included in the CCMP header. The PN and portions of the MAC header are used to generate a nonce that is in turn used by the CCM encryption process.

Four-Way Handshake

The four-way handshake is the method used to derive the encryption keys to encrypt wireless data frames. [Figure 4-7](#) graphically represents the frame exchanges used to generate the encryption keys. These keys are referred to as temporal keys.

The encryption keys are derived from the PMK that is mutually derived during the EAP authentication. This PMK is sent to the authenticator in the EAP success message, but is not forwarded to the supplicant because the supplicant has derived its own copy of the PMK.

1. The authenticator sends an EAPOL-Key frame containing an authenticator nonce (ANonce), which is a random number generated by the authenticator.
 - a. The supplicant derives a PTK from the ANonce and supplicant nonce (SNonce), which is a random number generated by the client/supplicant.
2. The supplicant sends an EAPOL-Key frame containing an SNonce, the RSN information element from the (re)association request frame, and an MIC.
 - a. The authenticator derives a PTK from the ANonce and SNonce and validates the MIC in the EAPOL-Key frame.
3. The authenticator sends an EAPOL-Key frame containing the ANonce, the RSN information element from its beacon or probe response messages; the MIC, determining whether to install the temporal keys; and the encapsulated group temporal key (GTK), the multicast encryption key.
4. The supplicant sends an EAPOL-Key frame to confirm that the temporal keys are installed.

Figure 4-7 Four-Way Handshake

Proactive Key Caching and CCKM

Proactive Key Caching (PKC) is an 802.11i extension that allows for the proactive caching (before the client roaming event) of the PMK that is derived during a client 802.1x/EAP authentication at the AP (see [Figure 4-8](#)). If a PMK (for a given WLAN client) is pre-cached at an AP, to which the client is about to roam, full 802.1x/EAP authentication is not required. Instead, the WLAN client can simply use the WPA four-way handshake process to securely derive a new session encryption key for communication with that AP.

The distribution of these cached PMKs to APs is greatly simplified in the Cisco Unified Wireless Network deployment. The PMK is simply cached in the controller(s) and made available to all APs that connect to it. The PMK is also shared with all other controllers that make up a mobility group with the anchor controller.

[illegible]

The state of the key cache for each WLAN client can be seen with the `show pmk-cache all` command. This identifies which clients are caching the keys, and which key caching mechanism is being used. The 802.11r workgroup is responsible for the standardization of an FSR mechanism for 802.11.

```

WLAN Identifier..... 1
Network Name (SSID)..... wpa2
...
Security
    802.11 Authentication:..... Open System
    Static WEP Keys..... Disabled
    802.1X..... Disabled
    Wi-Fi Protected Access (WPA/WPA2)..... Enabled
        WPA (SSN IE)..... Disabled
        WPA2 (RSN IE)..... Enabled
            TKIP Cipher..... Disabled
            AES Cipher..... Enabled
    Auth Key Management
        802.1x..... Enabled
        PSK..... Disabled
        CCKM..... Enabled

```

Type	Station	Entry Lifetime	VLAN Override	IP Override
CCKM	00:12:f0:7c:a3:47	43150		0.0.0.0
RSN	00:13:ce:89:da:8f	42000		0.0.0.0

Cisco Unified Wireless Network Architecture

Figure 4-9 shows a high level topology of the Cisco Unified Wireless Network architecture that includes CAPWAP APs, mesh CAPWAPs, the management system (WCS/NCS/PI), and the wireless LAN controller (WLC).

The Cisco Access Control Server (ACS) or the Identity Services Engine (ISE) and their AAA features complete the solution by providing RADIUS services in support of wireless user authentication and authorization.

Figure 4-9 Cisco Unified Wireless Network Architecture

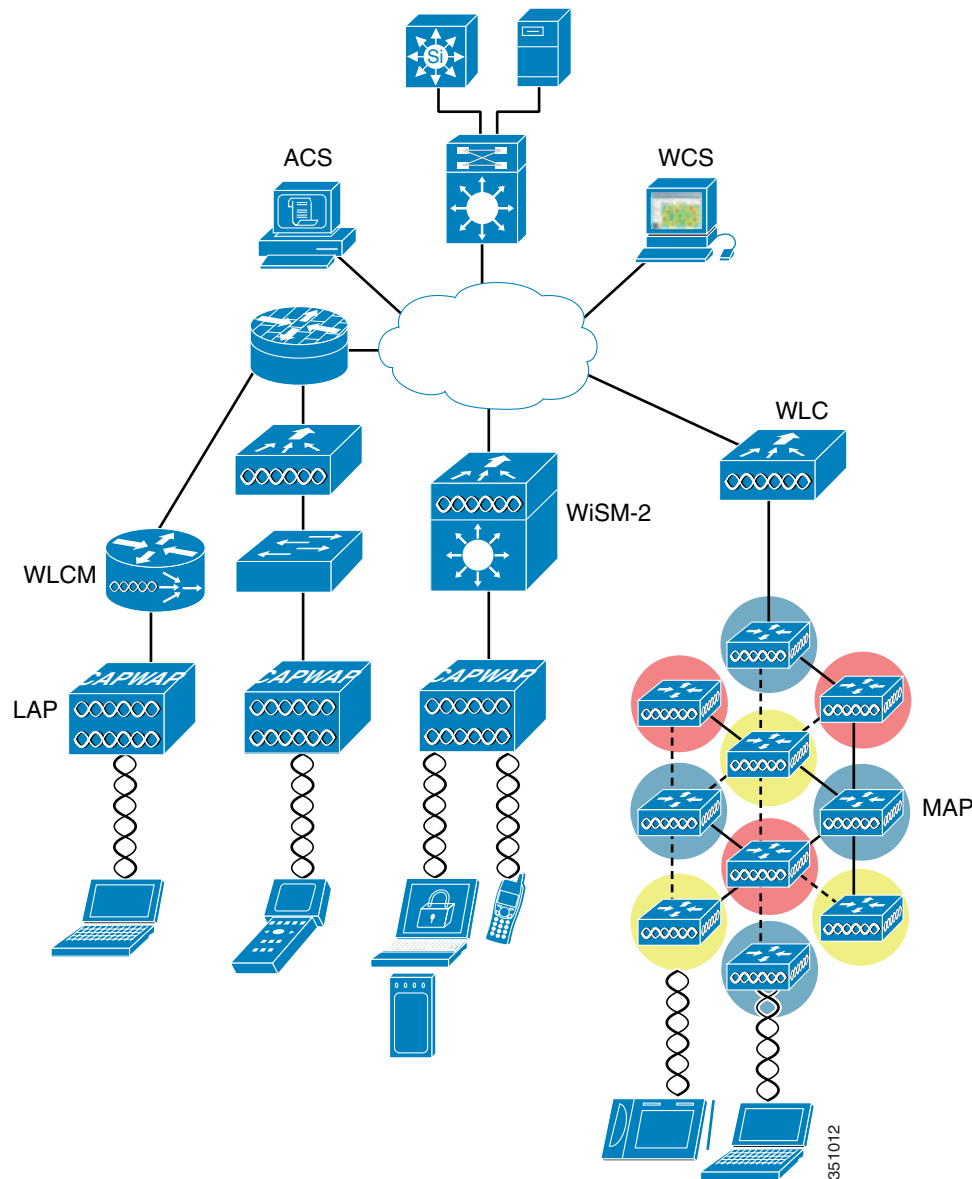
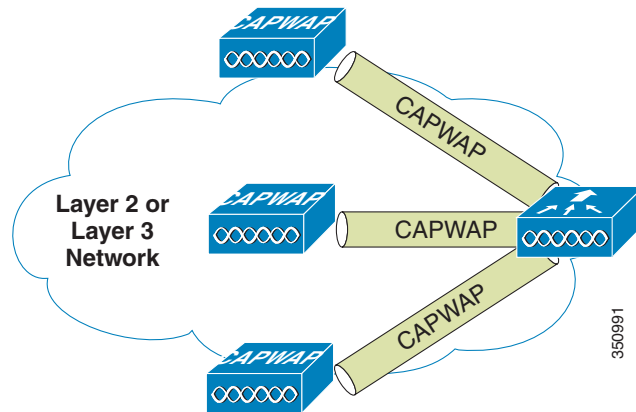


Figure 4-10 illustrates one of the primary features of the architecture: how APs use the CAPWAP protocol to communicate with and tunnel traffic to a WLC.

Figure 4-10 CAPWAP APs and WLC Connection



CAPWAP has three primary functions:

- Control and management of the AP
- Tunneling of WLAN client traffic to the WLC
- Collection of 802.11 data for the management of the Cisco Unified Wireless Network

CAPWAP Features

Control and Provisioning of Wireless Access Points Protocol (CAPWAP) is an update to Lightweight Access Point Protocol (LWAPP). CAPWAP is a standard, interoperable protocol that enables the WLC to manage a collection of APs. Features of CAPWAP include:

- An upgrade path from Cisco LWAPP products to next-generation Cisco products that use CAPWAP
- The ability to manage RFID readers and similar devices
- Controllers to interoperate with third-party access points

LWAPP-enabled APs can discover and join a CAPWAP controller; conversion to a CAPWAP controller is seamless. For example, the WLC discovery process and the firmware downloading process are the same with CAPWAP and LWAPP.

Important Points to Remember

- If your firewall is currently configured to allow traffic only from APs that use LWAPP, you must change the rules of the firewall to allow traffic from APs that use CAPWAP.
- Make sure that the CAPWAP UDP ports 5246 and 5247 (similar to the LWAPP UDP ports 12222 and 12223) are enabled and are not blocked by an intermediate device that could prevent an AP from joining the controller.
- If access control lists (ACLs) are in the control path between the controller and its APs, you need to open new protocol ports to prevent access points from being stranded.

APs use a random UDP source port to reach the destination ports on the controller. If you have a new out-of-the-box AP, it might try to contact the controller using LWAPP before it downloads the CAPWAP image from the controller. Once the AP downloads the CAPWAP image from the controller, it uses only CAPWAP to communicate with the controller.

**Note**

After 60 seconds of trying to join a controller with CAPWAP, the AP falls back to using LWAPP. If it cannot find a controller using LWAPP within 60 seconds, it tries again to join a controller using CAPWAP. The AP repeats this cycle of switching from CAPWAP to LWAPP and back again every 60 seconds until it joins a controller.

Cisco Unified Wireless Network Security Features

The native 802.11 security features combined with the physical security and ease of deployment of the CAPWAP architecture serves to improve the overall security of WLAN deployments. In addition to the inherent security benefits offered by the CAPWAP protocol, the Cisco Unified Wireless Network solution also includes the following additional security features:

- Enhanced WLAN security options
- ACL and firewall features
- Dynamic Host Configuration Protocol (DHCP) and Address Resolution Protocol (ARP) protection
- Peer-to-peer blocking
- Wireless intrusion protection system (wIPS)
 - Client exclusion
 - Rogue AP detection
- Management frame protection
- Dynamic RF management
- Architecture integration
- IDS integration

Enhanced WLAN Security Options

The Cisco Unified Wireless Network solution supports multiple concurrent WLAN security options. For example, multiple WLANs can be created on a WLC, each with its own WLAN security settings that can range from an open guest WLAN network and WEP networks for legacy platforms to combinations of WPA and/or WPA2 security configurations.

Each WLAN SSID can be mapped to either the same or different dot1q interface on the WLC, or Ethernet over IP (EoIP) tunneled to a different controller through a mobility anchor (Auto Anchor Mobility) connection.

If a WLAN client authenticates via 802.1x, a dot1q VLAN assignment can be controlled by way of RADIUS attributes passed to the WLC upon successful authentication.


Figure 4-11 and Figure 4-12 show a subset of the Unified Wireless Network WLAN configuration screen. The following three main configuration items appear on these screens:

- The WLAN SSID
- The WLC interface to which the WLAN is mapped
- The security method (Figure 4-12)

Figure 4-11 *WLANs General Tab*



Figure 4-12 *WLANs Layer 2 Security Tab*



Local EAP Authentication

The WLC software provides local EAP authentication capabilities that can be used when an external RADIUS server is not available or becomes unavailable. The delay before switching to local authentication is configurable, as illustrated in Figure 4-13. When RADIUS server availability is restored, the WLC automatically switches back from local authentication to RADIUS server authentication.

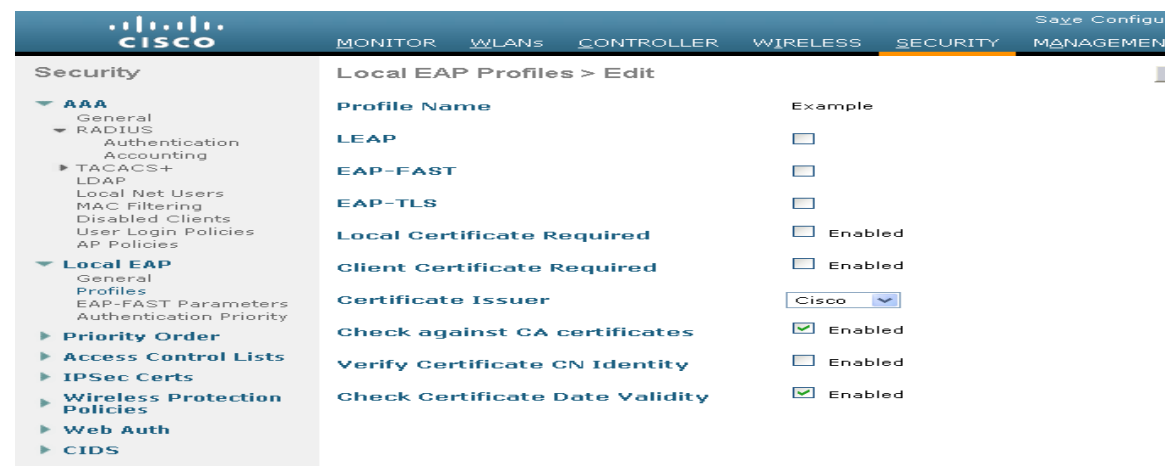
Figure 4-13 Local Authentication Timeout



The EAP types supported locally on the WLC are LEAP, EAP-FAST, EAP-TLS, and PEAP.


Figure 4-14 displays the window where you select the local EAP profiles.

Figure 4-14 Local EAP Profiles



WLC can use its local database for authentication data, and it can also access an LDAP directory to provide data for EAP-FAST or EAP-TLS authentication. The user credential database priority (LDAP versus Local) is configurable, as shown in [Figure 4-15](#).

Figure 4-15 Local EAP Priority



ACL and Firewall Features

The WLC allows access control lists (ACLs) to be defined for any interface configured on the WLC, as well as ACLs to be defined for the CPU of the WLC itself. These ACLs can be used to enforce policy on specific WLANs to limit access to particular addresses and/or protocols, as well as to provide additional protection to the WLC itself.

Interface ACLs act on WLAN client traffic in and out of the interfaces to which the ACLs are applied. CPU ACLs are independent of interfaces on the WLC, and are applied to all traffic to and from the WLC system.

[Figure 4-16](#) displays the ACL Configuration page. The ACL can specify source and destination address ranges, protocols, source and destination ports, DSCP, and direction in which the ACL is to be applied. An ACL can be created out of a sequence of various rules.

Figure 4-16 ACL Configuration Page

The screenshot displays the Cisco configuration interface for Access Control Lists (ACLs). On the left, a navigation tree under 'Security' includes sections like AAA, Local EAP, Priority Order, and 'Access Control Lists' (which is highlighted with a red circle). The main area is titled 'Access Control Lists > Rules > New'. It contains several configuration fields: 'Sequence' (text input with '10'), 'Source' (dropdown menu set to 'Any'), 'Destination' (dropdown menu set to 'Any'), 'Protocol' (dropdown menu set to 'UDP'), 'Source Port' (dropdown menu set to 'Any'), 'Destination Port' (dropdown menu set to 'Any'), 'DSCP' (dropdown menu set to 'Any'), 'Direction' (dropdown menu set to 'Any'), and 'Action' (dropdown menu set to 'Deny').

DHCP and ARP Protection

The WLC acts as a relay agent for WLAN client DHCP requests. In doing so, the WLC performs a number of checks to protect the DHCP infrastructure. The primary check is to verify that the MAC address included in the DHCP request matches the MAC address of the WLAN client sending the request. This protects against DHCP exhaustion attacks, by restricting a WLAN client to one DHCP request (IP address) for its own interface. The WLC by default does not forward broadcast messages from WLAN clients back out onto the WLAN, which prevents a WLAN client from acting as a DHCP server and spoofing incorrect DHCP information.

The WLC acts as an ARP proxy for WLAN clients by maintaining the MAC address-IP address associations. This allows the WLC to block duplicate IP address and ARP spoofing attacks. The WLC does not allow direct ARP communication between WLAN clients. This also prevents ARP spoofing attacks directed at WLAN client devices.

Peer-to-Peer Blocking

The WLC can be configured to block communication between clients on the same WLAN. This prevents potential attacks between clients on the same subnet by forcing communication through the router.

[Figure 4-17](#) is the configuration screen for peer-to-peer blocking on the WLC. Note that this is a global setting on the WLC and applies to all WLANs configured on the WLC.

Figure 4-17 *Peer-to-Peer Blocking*

Controller General

802.3x Flow Control Mode: Disabled

LWAPP Transport Mode: Layer 3 (Current Operating Mode is Layer3)

LAG Mode on next reboot: Enabled (LAG Mode is currently enabled).

Ethernet Multicast Mode: Disabled

Broadcast Forwarding: Disabled

Aggressive Load Balancing: Disabled

Peer to Peer Blocking Mode: Disabled

Over The Air Provisioning of AP: Enabled

AP Fallback: Enabled

Apple Talk Bridging: Disabled

Fast SSID change: Disabled

Default Mobility Domain Name: SRND

RF-Network Name: srnd

Wireless IDS

The WLC performs WLAN IDS analysis using information obtained from all of the connected APs, and reports detected attacks to WLC as well to the WCS. The Wireless IDS analysis is complementary to any analysis that can otherwise be performed by a wired network IDS system. The embedded Wireless IDS capability of the WLC analyzes 802.11 and WLC-specific information that is not otherwise visible or available to a wired network IDS system.

The wireless IDS signature files used by the WLC are included in WLC software releases; however, they can be updated independently using a separate signature file. Custom signatures are displayed in the Custom Signatures window.

Figure 4-18 is the Standard Signatures window in the WLC.

Figure 4-18 *Standard WLAN IDS Signatures*

Security Standard Signatures

Global Settings

Enable check for all Standard and Custom Signatures: ☒


Signatures

Precedence	Name	Frame Type	Action	State	Description
1	Bcast deauth	Managemen	Report	Enabled	Broadcast Deauth
2	NULL probe resp 1	Managemen	Report	Enabled	NULL Probe Res
3	NULL probe resp 2	Managemen	Report	Enabled	NULL Probe Res
4	Assoc flood	Managemen	Report	Enabled	Association Rec
5	Reassoc flood	Managemen	Report	Enabled	Reassociation F
6	Broadcast Probe floo	Managemen	Report	Enabled	Broadcast Prob
7	Disassoc flood	Managemen	Report	Enabled	Disassociation f
8	Deauth flood	Managemen	Report	Enabled	Deauthentication
9	Res mgmt 6 & 7	Managemen	Report	Enabled	Reserved man
10	Res mgmt D	Managemen	Report	Enabled	Reserved man
11	Res mgmt E & F	Managemen	Report	Enabled	Reserved man
12	EAPOL flood	Data	Report	Enabled	EAPOL Flood At
13	NetStumbler 3.2.0	Data	Report	Enabled	NetStumbler 3.2.0

Cisco Adaptive Wireless Intrusion Prevention System

The Cisco Adaptive Wireless Intrusion Prevention System (wIPS) with ELM feature allows you to provide comprehensive security protection to your deployed APs without the need of a dedicated monitor mode or an overlay network (Figure 4-19). APs must provide protection from unauthorized security access, penetration, and attacks. Cisco wIPS with the ELM feature enabled on your network APs effectively eases the implementation of wireless security.

Figure 4-19 AP Deployment with Enhanced Local Mode (ELM)



The wIPS Communication Protocols in Figure 4-19 are:

- **CAPWAP**—This is the successor to Lightweight Access Point Protocol (LWAPP) and is utilized for communication between ELM APs and WLC. It provides a bi-directional tunnel in which alarm information is shuttled from the WLC to wIPS and other Cisco Prime Infrastructure management system configuration information is pushed to the AP.



Note

The Cisco Prime Infrastructure management system was formerly known as Wireless Control System (WCS), which evolved to Network Control System (NCS). For clarity, all three are referred to as *management system (WCS/NCS/PI)*.


- Network Mobility Services Protocol (NMSP)—This encrypted protocol communicates between the WLC and the management system (WCS/NCS/PI). In a wIPS deployment, this protocol provides a pathway for alarm information to aggregate from WLC to wIPS (and other services running) and for wIPS configuration information to be pushed to the controller.
- SOAP/XML (Simple Object Access Protocol)—The method of communication to the management system (WCS/NCS/PI). This protocol is used to distribute configuration parameters to the wIPS and other services running on the Mobility Service Engine (MSE).
- SNMP (Simple Network Management Protocol)—Used to forward wIPS alarm information from the MSE to the management system (WCS/NCS/PI). It is also used to communicate rogue AP information from the WLC to the management system (WCS/NCS/PI).

Dedicated Monitor Mode versus ELM

Figure 4-20 illustrates a contrast between the standard deployments of wIPS monitor mode and APs with the ELM feature. The typical coverage range for both modes suggests:

- Dedicated wIPS monitor mode APs (shown in red in Figure 4-20) typically covers 15,000 to 35,000 square feet
- APs with the ELM feature (shown in yellow in Figure 4-20) typically cover from 3,000 to 5,000 square feet

Figure 4-20 Monitor Mode versus ELM



In the traditional wIPS deployment, a recommended ratio is 1 monitor mode AP to every 5 local mode APs (ratio can vary based on network design and expert guidance for best coverage). With ELM, you simply enable the ELM feature for all of the APs, effectively adding monitor mode wIPS operations to local data-serving mode AP while still maintaining performance.

On-Channel and Off-Channel Performance

When an AP visits a channel, the time the AP *stays* on that channel, to detect and classify an attack, is known as the *dwel* time. ELM primary feature operates effectively for on-channel attacks, without any compromise to the performance on data, voice and video clients, and services. In contrast, the local mode varies off-channel scanning providing minimal dwell time to detect and classify an attack.

For example, due to radio resource management (RRM), when voice clients are associated to an AP scanning is deferred until the voice client is disassociated in order to ensure service is not affected. In this example, ELM detection during off-channel is considered best effort. Neighboring ELM APs operating on all/country/DCA channels increases effectiveness, hence the recommendation for enabling ELM on every local mode AP for maximum coverage protection. If your requirement is for dedicated scanning on all channels full-time, then Cisco recommends deploying monitor mode APs.

Generally, the differences between local mode and monitor mode APs are:

- Local Mode AP—Serves WLAN clients with time slicing off-channel scanning, listens for 50 ms on each channel, and features configurable scanning for all/country/DCA channels.
- Monitor Mode AP—Does not serve WLAN clients, dedicated to scanning only, listens for 1.2 sec on each channel, and scans all channels.

ELM Across WAN Links

Cisco has optimized features in challenging topologies, such as deploying ELM APs across low bandwidth WAN links. The ELM feature involves pre-processing to determine attack signatures at the AP and is optimized to work over slower links. The Cisco recommended best practice is to test and measure the baseline to validate performance with ELM over WAN.

CleanAir Integration

Cisco CleanAir technology is a spectrum-aware, self-healing, and self-optimizing wireless network that mitigates the impact of wireless interference and offers performance protection for 802.11n networks.

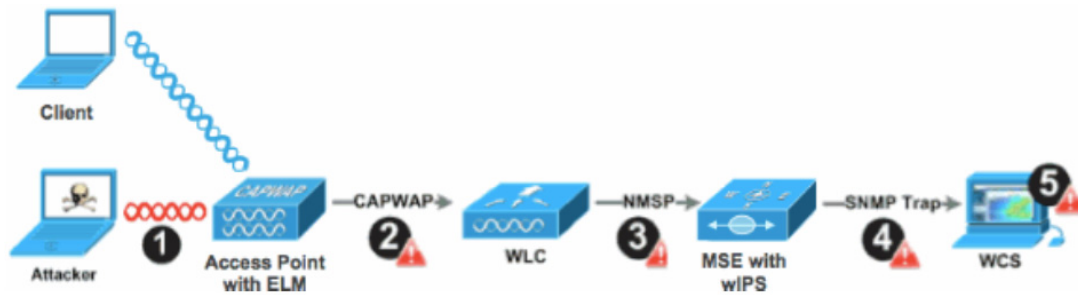
The ELM feature compliments CleanAir operations with similar performance and benefits as monitor mode AP deployments, including these existing CleanAir spectrum-aware benefits:

- Dedicated silicon-level RF intelligence
- Spectrum-aware, self-healing, and self-optimizing
- Non-standard channel threat and interference detection and mitigation
- Non-Wi-Fi detection such as Bluetooth, microwave, cordless phones, and so forth
- Detect and locate RF layer DOS attacks such as RF jammers

ELM wIPS Alarm Flow

Attacks are only relevant when they occur on *trusted* APs. The ELM APs will detect an attack, then communicate, correlate, and report to the management system (WCS/NCS/PI), as shown in [Figure 4-21](#). Generally, the alarm flow process is:


1. Attack is launched against a *trusted* AP
2. Detection on the AP with ELM feature communicates through CAPWAP to WLC
3. Passed transparently to MSE via NMSP
4. Log into wIPS database on MSE and send to the management system (WCS/NCS/PI) by way of an SNMP trap
5. Display at the management system (WCS/NCS/PI)

Figure 4-21 Threat Detection Alarm Flow

Client Exclusion

In addition to Wireless IDS, the WLC is able to take additional steps to protect the WLAN infrastructure and WLAN clients. The WLC is able to implement policies that exclude WLAN clients whose behavior is considered threatening or inappropriate. Figure 4-22 shows the Exclusion Policies window, containing the following currently supported client exclusion policies:

- Excessive 802.11 association failures—Possible faulty client or DoS attack
- Excessive 802.11 authentication failures—Possible faulty client or DoS attack
- Excessive 802.1X authentication failures—Possible faulty client or DoS attack
- IP theft or IP reuse—Possible faulty client or DoS attack
- Excessive web authentication failures—Possible DoS or password-cracking attack

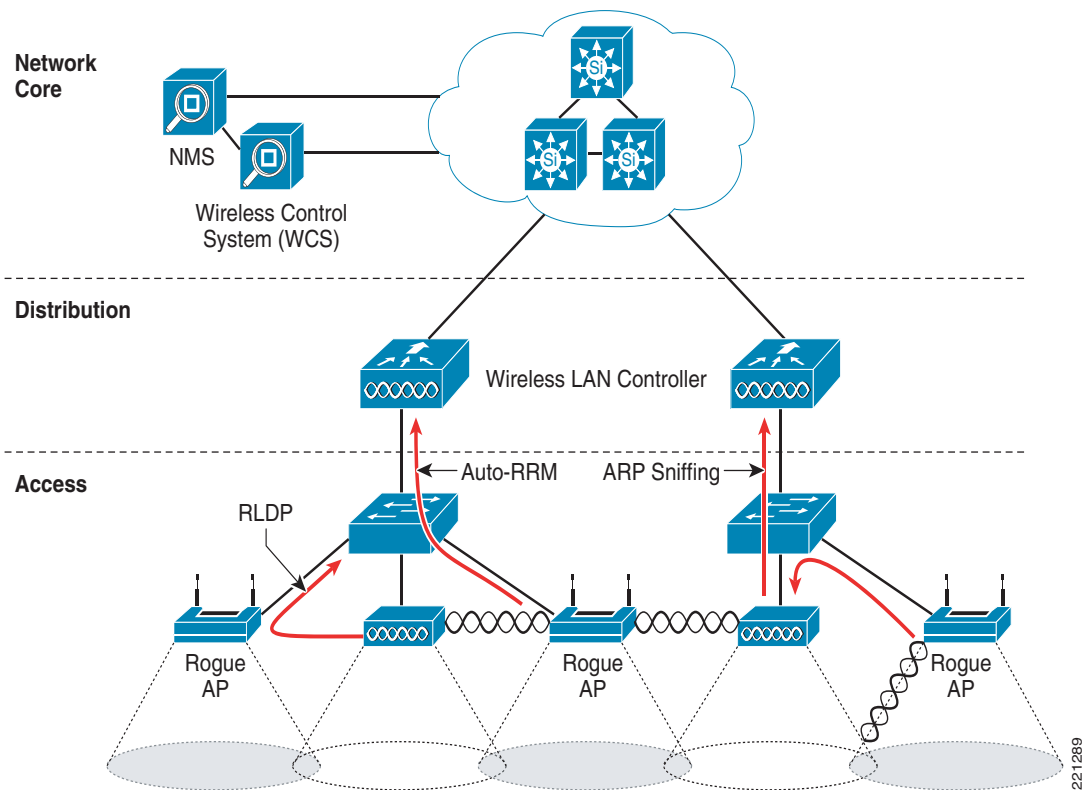
Figure 4-22 Client Exclusion Policies

Rogue AP

The Cisco Unified Wireless Networking solution, as shown in Figure 4-23, provides a complete solution for rogue APs. This solution provides:

- Air/RF detection—Detection of rogue devices by observing/sniffing beacons and 802.11 probe responses.
- Rogue AP location—Use of the detected RF characteristics and known properties of the managed RF network to locate the rogue device.
- Wire detection—A mechanism for tracking/correlating the rogue device to the wired network.
- Rogue AP isolation —A mechanism to prevent client connection to a rogue AP.

Figure 4-23 Unified Wireless Network Rogue AP Detection



Air/RF Detection

The two AP RF detection deployment models are:

- Standard AP deployment
- Monitor mode AP deployment

Both deployment models support RF detection and are not limited to rogue APs, but can also capture information upon detection of ad-hoc clients and rogue clients (the users of rogue APs). An AP that is configured for monitor mode is dedicated to scanning the RF channels and does not support client association or data transmission.

When searching for rogue APs, an AP goes off channel for 50 ms to listen for rogue clients, and to monitor noise and channel interference. The channels scanned are configured in the global WLAN network parameters for 802.11a and 802.11b/g.

Any detected prospective rogue client(s) and/or access points are sent to the controller to gather the following information:

- Rogue AP MAC address
- Rogue AP name
- Rogue connected client(s) MAC address
- Whether the frames are protected with WPA, WEP and WEP2
- The preamble
- Signal-to-noise ratio (SNR)
- Received signal strength indication (RSSI)
- Switchport tracing

The prospective rogue client/AP is not labeled a rogue until the WLC receives another report from a trusted AP or until the completion of a second detection cycle. The trusted AP moves to the same channel, as the prospective rogue, to monitor for rogue client/AP, noise, or interference. If the same client/AP is detected a second time, they are then labeled as *rogue* on the WLC.

Once labeled as a rogue, the WLC determines if this rogue is attached to the local network or is simply a neighboring AP. In either case, an AP that is not part of the managed Cisco Unified Wireless Network is considered a rogue.

In monitor mode, the trusted AP does not carry user traffic; it is dedicated to scanning channels. This mode of deployment is most common when a customer does not want to support WLAN services in a particular area, but wants to monitor that area for rogue APs and rogue clients.

Location

The location features of Cisco Prime Infrastructure can be used to provide a floor plan indicating the approximate location of a rogue AP. The floor plan displays the location of all legitimate APs, and highlights the location of a rogue AP with the skull-and-crossbones icon. For additional information on the Cisco Unified Wireless Network location features, see:

<http://www.cisco.com/en/US/products/ps6386/index.html>.

Wire Detection

Situations can exist where the Cisco Prime Infrastructure rogue location feature is not effective, such as in branch offices with only a few APs or where floor plan information might not be available. In these cases, the Cisco Unified Wireless Network solution offers two wire-based detection options:

- Rogue detector AP
- Rogue Location Discovery Protocol (RLDP)

If an AP is configured as a rogue detector, its radio is turned off and its role is to listen on the wired network for MAC addresses of clients associated to rogue APs; that is, *rogue clients*. The rogue detector listens for ARP packets that include rogue client MAC addresses. When it detects one of these ARPs, it reports this to the WLC, providing verification that the rogue AP is attached to the same network as the Cisco Unified Wireless Network.

To maximize the likelihood of capturing ARP information, the rogue AP detector is connected to all available broadcast domains using a Switched Port Analyzer (SPAN) port. Multiple rogue AP detector APs can be deployed to capture the various aggregated broadcast domains that exist on a typical network.

If a rogue client resides behind a wireless router (a common home WLAN device), its ARP requests are not seen on the wired network, so an alternative to the rogue detector AP method is needed. Additionally, rogue detector APs might not be practical for some deployments because of the large number of broadcast domains to be monitored (such as in the main campus network).

The RLDLP option can aid in these situations. In this case, a standard AP, upon detecting a rogue AP, can attempt to associate with the rogue AP as a client and send a test packet to the controller, which requires the AP to stop behaving as a standard AP and temporarily go into client mode. This action confirms that the rogue AP in question is actually on the network, and provides IP address information that indicates its logical location in the network. Given the difficulties in deriving location information in branch offices coupled with the likelihood of a rogue being located in multi-tenant buildings, rogue AP detector and RLDLP are useful tools that augment location-based rogue AP detection.

Switch Port Tracing

Cisco Prime Infrastructure provides rogue access point detection by retrieving information from the controller. The rogue access point table is populated with any detected BSSID addresses from any frames that are not present in the *neighbor list*. A neighbor list contains the known BSSID addresses of validated APs or *neighbors*. At the end of a specified interval, the contents of the rogue table are sent to the controller in a CAPWAP Rogue AP Report message. With this method, the Cisco Prime Infrastructure simply gathers the information received from controllers. Additionally, you can also incorporate auto or manual switch port tracing (SPT) of wired rogue access point switch ports. The auto SPT is preferable for a large wireless network.

Auto SPT launches automatically when a rogue AP is reported to the Cisco Prime Infrastructure. The auto SPT provides a quicker scan based on the wired location association of the rogue AP. The Cisco Prime Infrastructure allows you to configure the criteria for auto SPT and auto containment so that you can run a trace and contain the detected rogue access points on the wire.

When the multiple controllers report that a rogue AP should be auto contained, the Cisco Prime Infrastructure finds the controller that reports the strongest RSSI and sends the containment request to the controller.

Rogue AP Containment

Rogue AP connected clients, or rogue ad-hoc connected clients, can be contained by sending 802.11 de-authentication packets from nearby APs. This should be done only after steps have been taken to ensure that the AP is truly a rogue AP, because it is illegal to do this to a legitimate AP in a neighboring WLAN. This is why Cisco removed the automatic rogue AP containment feature from the solution.

To determine whether rogue AP clients are also clients on the enterprise WLAN, the client MAC address can be compared with MAC addresses collected by the AAA during 802.1X authentication. This allows for the identification of potential WLAN clients that might have been compromised or users who are not following security policies.


Management Frame Protection

One of the challenges in 802.11 has been that management frames are sent in the clear with no encryption or message integrity checking and are therefore vulnerable to spoofing attacks. WLAN management frame spoofing can be used to attack a WLAN network. To address this, Cisco created a digital signature

mechanism to insert a message integrity check (MIC) into 802.11 management frames. This allows legitimate members of a WLAN deployment to be identified, as well as being able to identify rogue infrastructure devices, and spoofed frames through their lack of valid MICs.

The MIC used in management frame protection (MFP) is not a simple CRC hashing of the message, but also includes a digital signature component. The MIC component of MFP ensures that a frame has not been tampered with, and the digital signature component ensures that the MIC could have only been produced by a valid member of the WLAN domain. The digital signature key used in MFP is shared among all controllers in a mobility group; different mobility groups have different keys allowing validation of all WLAN management frames processed, by the WLCs, in that mobility group (Figure 4-24).

Figure 4-24 Management Frame Protection

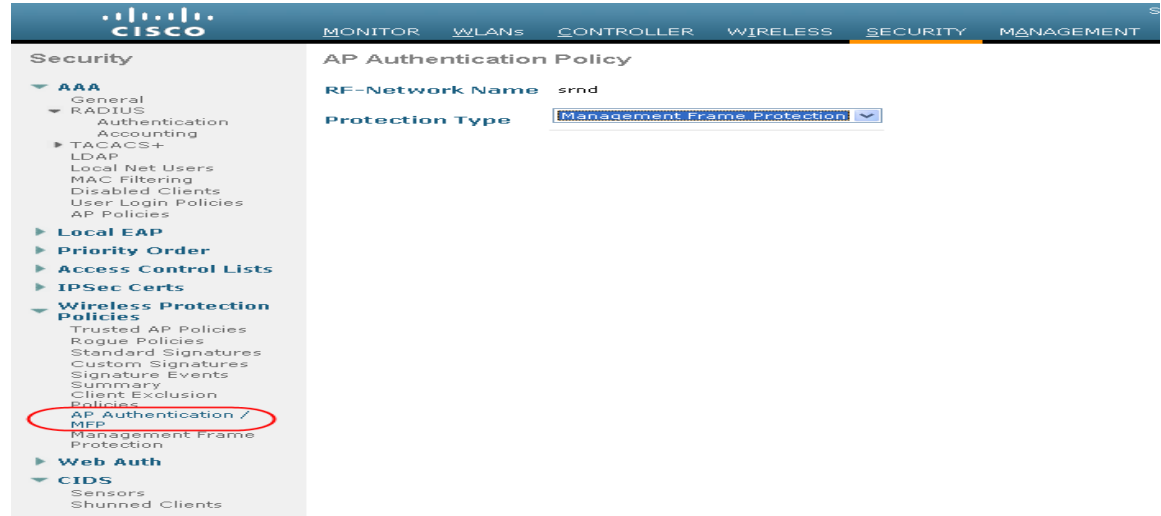
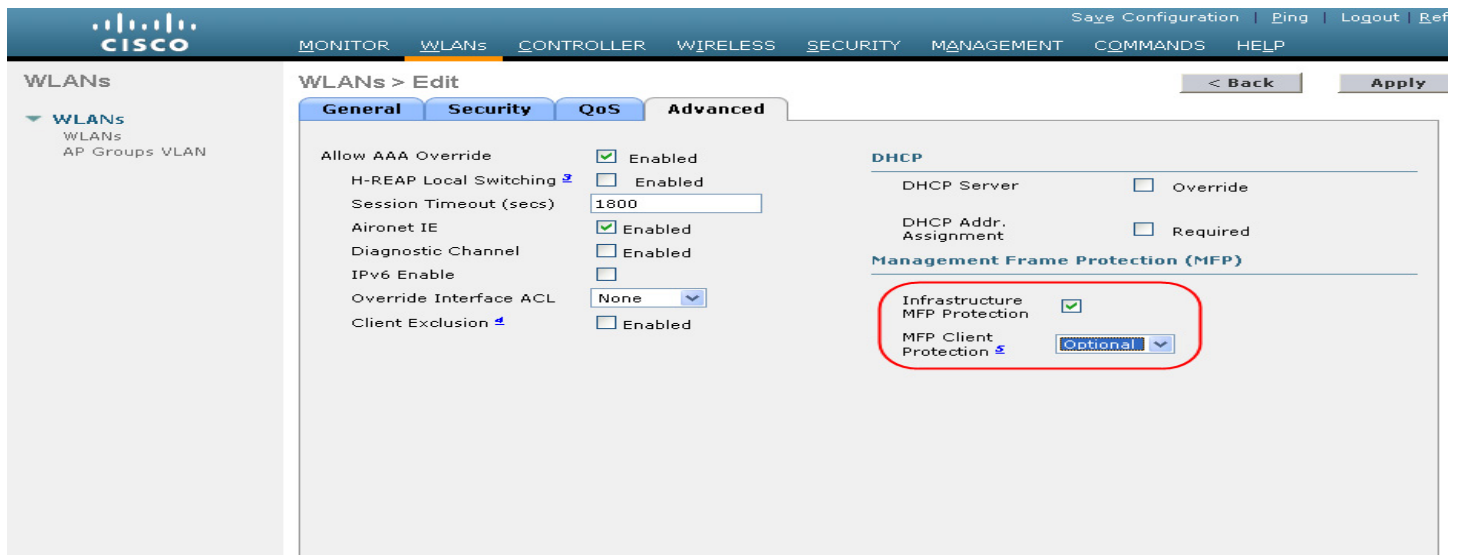


Both infrastructure-side and client MFP are currently possible, but client MFP requires Cisco Compatible Extensions v5 WLAN clients to *learn* the mobility group MFP key before they can detect and reject invalid frames.

MFP provides the following benefits:

- Authenticates 802.11 management frames generated by the WLAN network infrastructure
- Allows detection of malicious rogues that spoof a valid AP MAC or SSID to avoid detection as a rogue AP, or as part of a man-in-the-middle attack
- Increases the effectiveness of the rogue AP and WLAN IDS signature detection of the solution
- Provides protection of client devices using Cisco Compatible Extensions v5
- Supported by standalone AP/WDS/WLSE in version 12.3(8)/v2.13

Two steps are required to enable MFP: enabling it under the Security tab on the WLC (Figure 4-25) and enabling it on the WLANs in the mobility group (Figure 4-26).

Figure 4-25 Enabling MFP on the Controller**Figure 4-26** Enabling MFP per WLAN

Client Management Frame Protection

Cisco Compatible Extensions v5 WLAN clients support MFP. This is enabled on a per-WLAN basis, as is shown in [Figure 4-26](#) above.

The method of providing MFP for WLAN clients is fundamentally the same as that used for APs, which is to use a MIC in the management frames. This allows trusted management frames to be identified by the client. MIC cryptographic keys are passed to the client during the WPA2 authentication process. Client MFP is available only for WPA2. If WPA and WPA2 clients share the same WLAN, client MFP must be set to “optional”.

Management System Security Features

Apart from providing location support for Rogue AP detection, the management system (WCS/NCS/PI) provides two additional Unified Wireless Network security features: WLC configuration verification management and an alarm and reporting interface.

Configuration Verification

The management system (WCS/NCS/PI) can provide on-demand or regularly-scheduled configuration audit reports, which compare the complete current running configuration of a WLC and its registered access points with that of a known valid configuration stored in the management system (WCS/NCS/PI) databases. Any exceptions between the current running configuration and the stored database configuration are noted and brought to the attention of the network administrator via screen reports (Figure 4-27).

Figure 4-27 Audit Report Example

171.71.128.75 > Audit Report

Device name	171.71.128.75	Time of Audit	1:00:23
Report ID	68	Synchronization Status	Different In WCS And Controller

Object name	802.11 171.71.128.75
Synchronization Status	Different In WCS And Controller

<

Attribute	Value In WCS	Value In Device
bridgingSharedSecretKey	*****	*****

Object name	Known Rogues 171.71.128.75 00:01:64:45:b9:b8
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:0e:37:bf
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:1f:93:f9
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:1f:94:15
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:40:4d
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:41:01
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f0
Synchronization Status	Not Present In Controller

Object name	Known Rogues 171.71.128.75 00:02:8a:5b:46:f1
Synchronization Status	Not Present In Controller

190735

Alarms and Reports

Apart from the alarms that can be generated directly from a WLC and sent to an enterprise network management system (NMS), the management system can also send alarm notifications. The primary difference between alarm notification methods, apart from the type of alarm sent by the various components, is that the WLC uses Simple Network Management Protocol (SNMP) traps to send alarms (which can be interpreted only by an NMS system), whereas the management system (WCS/NCS/PI) uses SMTP e-mail to send an alarm message to an administrator.

The management system (WCS/NCS/PI) provides both real-time and scheduled reports, and can export or e-mail reports. The management system (WCS/NCS/PI) provides reports on:

- Access points
- Audits
- Clients
- Inventory
- Mesh
- Performance
- Security

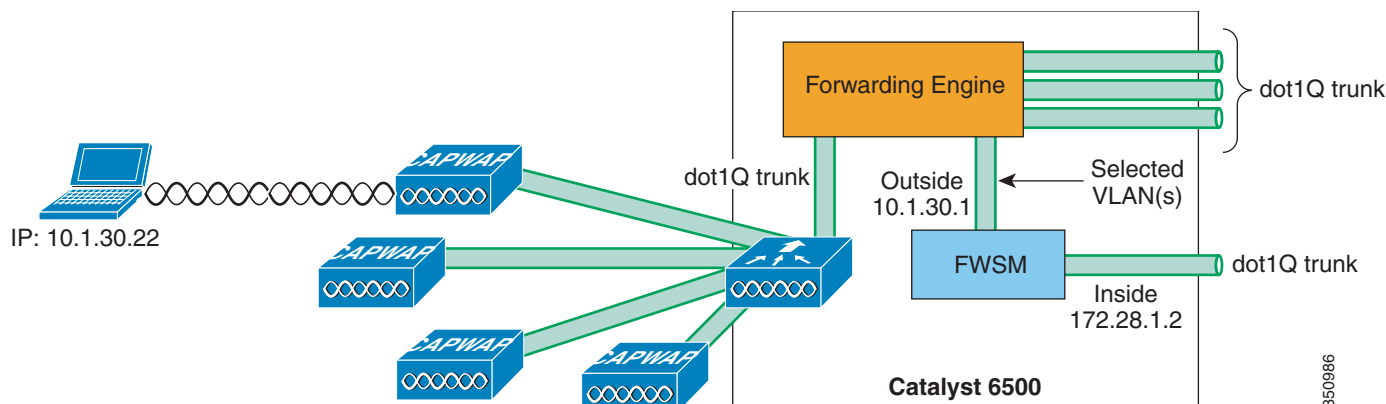
Architecture Integration

Cisco provides a wide variety of security services that are either integrated into Cisco IOS, integrated into service/network modules, offered as standalone appliances, or as software.

The Cisco Unified Wireless Network architecture eases the integration of these security services into the solution because it provides a Layer 2 connection between the WLAN clients and the upstream wired network. This means that appliances or modules that operate by being “inline” with client traffic can be easily inserted between WLAN clients and the wired network. For example, an older WLSM-based deployment requires the implementation of VRF-Lite on the Cisco 6500 to enable WLAN client traffic to flow through a Cisco Firewall Service Module (FWSM); whereas in a Cisco Unified WLAN deployment, a WiSM can simply map the (WLAN) client VLAN directly to the FWSM. The only WLAN controllers in the Cisco Wireless portfolio that cannot directly map WLAN traffic to a physical/logical interface at Layer 2 are ISR-based WLC modules. An ISR WLAN module does have access to all the IOS and IPS features available on the ISR, but IP traffic from the WLAN clients must be directed in and out of specific ISR service module interfaces using IOS VRF features on the router.

Figure 4-28 shows an example of architectural integration between a WiSM and the FWSM module. In this example, the WLAN client is on the same subnet as the outside firewall interface. No routing policy or VRF configuration is required to ensure that WLAN client traffic in both directions goes through the firewall.

A Cisco Network Admission Control (NAC) Appliance (formerly Cisco Clean Access) can be implemented in combination with a WLAN deployment to ensure that end devices connecting to the network meet enterprise policies for compliance with latest security software requirements and operating system patches. Like the FWSM module discussed above, the Cisco NAC Appliance can also be integrated into a Cisco Unified Wireless Network architecture at Layer 2, thereby permitting strict control over which wireless user VLANs are subject to NAC policy enforcement.

Figure 4-28 Firewall Module Integration Example

In addition to ease of integration at the network layer, the Cisco Unified Wireless Network solution provides integration with Cisco IDS deployments, allowing clients blocked by the Cisco IDS to be excluded from the Cisco Unified Wireless Network.

Cisco Integrated Security Features

Cisco Integrated Security Features (CISF) are available on Cisco Catalyst switches, and help mitigate against a variety of attacks that a malicious user might launch after gaining wireless access to the network. This section describes these attacks, how a WLC protects against these attacks, and how CISF, when enabled on the access switch, can help protect the network.



Note

This section describes only the attacks that CISF can help prevent when enabled on access switches, and is not meant to be a comprehensive analysis of all the attacks that are possible on wireless networks.

Types of Attacks

Attacks can occur against either wired or wireless networks. However, a wireless network connection allows an attacker to craft an attack without physical connectivity to the network. The WLC and CISF include features that are specifically designed to prevent such attacks, including the following:

- MAC flooding attacks
- DHCP rogue server attacks
- DHCP exhaustion attacks
 - ARP spoofing attacks
 - IP spoofing attacks

MAC Flooding Attack

MAC flooding attacks are attempts to fill the content-addressable memory (CAM) table of a switch, and thus force the switch to start flooding LAN traffic. These attacks are performed with tools such as macof (part of the dsniff package), which generates a flood of frames with random MAC and IP source and destination addresses.

The Layer 2 learning mechanism of an Ethernet switch is based on the source MAC addresses of packets. For each new source MAC address received on a port, the switch creates a CAM table entry for that port and for the VLAN to which the port belongs. The *macof* utility typically fills the CAM table in less than ten seconds, given the finite memory available to store these entries on the switch. CAM tables are limited in size. If enough entries are entered into the CAM table before other entries expire, the CAM table fills up to the point that no new entries can be accepted.

When the CAM table of a switch is filled, it then floods all of its ports with incoming traffic because it cannot find the port number for a particular MAC address in the CAM table. The switch, in essence, acts like a hub to the detriment of performance and security. The overflow floods traffic within the local VLAN, so the intruder sees traffic within the VLAN to which he or she is connected.

At Layer 3, the random IP destinations targeted by *macof* also use the multicast address space. Thus, the distribution layer switches that have multicast turned on experience high CPU usage levels as the protocol independent multicast (PIM) process attempts to handle the false routes.

DHCP Rogue Server Attack

The DHCP rogue server event can be the result of a purposeful attack, or a user might have accidentally brought up a DHCP server on a network segment and begun to inadvertently issue IP addresses. An intruder can bring up a DHCP server and offer IP addresses representing a DNS server or default gateway that redirects unsuspecting user traffic to a computer under the control of the intruder.

DHCP Starvation Attack

DHCP starvation attacks are designed to deplete all of the addresses within the DHCP scope on a particular segment. Subsequently, a legitimate user is denied an IP address requested by way of DHCP and thus is not able to access the network. Gobbler is a public domain hacking tool that performs automated DHCP starvation attacks. DHCP starvation can be purely a DoS mechanism or can be used in conjunction with a malicious rogue server attack to redirect traffic to a malicious computer ready to intercept traffic.

ARP Spoofing-based Man-In-the-Middle Attack

A man-in-the-middle (MIM) attack is a network security breach in which a malicious user intercepts (and possibly alters) data traveling along a network. One MIM attack uses ARP spoofing, in which a gratuitous Address Resolution Protocol (ARP) request is used to misdirect traffic to a malicious computer such that the computer becomes the “man in the middle” of IP sessions on a particular LAN segment. The hacking tools *ettercap*, *dsniff*, and *arpspoof* can be used to perform ARP spoofing. Ettercap in particular provides a sophisticated user interface that displays all of the stations on a particular LAN segment and includes built-in intelligent packet capturing to capture passwords on a variety of IP session types.

IP Spoofing Attack

IP spoofing attacks spoof the IP address of another user to perform DoS attacks. For example, an attacker can ping a third-party system while sourcing the IP address of the second party under attack. The ping response is directed to the second party from the third-party system.

CISF for Wireless Deployment Topologies

This section describes the various Cisco Unified Wireless Network deployment topologies. The following section describes how the WLC or CISF features defend against wireless attacks.


CISF is currently available only on the access switch, not directly on the access point (AP); thus, the benefits of these features are available only if the traffic from the wireless attacker goes through the switch.

The definition of an access switch is slightly different in the Unified Wireless Network solution, because three locations can be considered an access switch:

- The point that a controller interface terminates on the network
- The point that a CAPWAP AP terminates on the network
- The point that a FlexConnect AP terminates on the network

These locations are illustrated in Figure 4-29.

Figure 4-29 Access Switches



The connections of interest to CISF are the controller switch and the FlexConnect switch. The AP switch is not discussed because WLAN traffic does not terminate on this switch, and the AP simply appears as a single device connected to that switch port, so from a security point of view it can be considered an access client.



Note

The primary difference between a CAPWAP AP and a standard client is that the differentiated services code point (DSCP) value of a CAPWAP AP is trusted.

The scope of the following topologies is limited to attacks between wireless users because it is assumed that wireless and wired users are supported on separate subnets (as recommended by Cisco best practices) and because any discussion of inter-subnet attacks is beyond the scope of this discussion.

The three following topologies are considered:

- Topology 1—Target is associated to the same AP to which the attacker is connected
- Topology 2—Target is associated to a different AP than the attacker
- Topology 3—Target is associated to a different AP than the attacker, and this AP is connected to a different controller

In first topology, both attacker and target are associated to the same AP, the traffic remains local to the FlexConnect or WLC, and CISF is not useful, but the Cisco Unified Wireless Networks native security address these issues. The second and third topologies are the ones in which CISF can be effective.

For an enterprise WLAN deployment requiring different levels of authorization, multiple VLANs per SSID are commonly used. This requires configuring an 802.1q trunk between the Fast Ethernet port on the FlexConnect AP or WLC, and the corresponding port on the access switch. With multiple VLANs defined, the administrator can keep the data traffic separated from the AP and WLC management traffic.

The company security policy is also likely to require having different types of authentication and encryptions for different types of users (open authentication and no encryption for guest access, dot1x authentication and strong encryption for employees, and so on). This is achieved by defining multiple SSIDs and VLANs on the FlexConnect AP or WLC.

Given the above, the configurations used in the example configurations assume a trunk connection between the WLC or FlexConnect AP and the access switch.

Using Port Security to Mitigate a MAC Flooding Attack

Port security sets a maximum number of MAC addresses allowed on a port. You can add addresses to the address table manually, dynamically, or by a combination of the two. Packets are dropped in hardware when the maximum number of MAC addresses in the address table is reached, and a station that does not have a MAC address in the address table attempts to send traffic.

Enabling port security on the access port of the switch stops a MAC flooding attack from occurring because it limits the MAC addresses allowed through that port. If the response to a violation is set to **shutdown**, the port goes to error-disable state. If the response is set to *restrict*, traffic with unknown source MAC addresses are dropped.


Port Security in a Wireless Network

It is not generally recommended to enable port security on a switch port connected to a FlexConnect AP or WLC. The use of port security implies knowing the exact number of MAC addresses that the switch learns and allows from that port; in the case of an FlexConnect AP or WLC, the various source MAC addresses that the switch learns usually correspond to wireless users. Setting port security on the switch port allows only a certain number of users on the wired network.

For example, a company might have a security policy that allows only a certain number of MACs to send traffic through the access point. In this case, a combination of MAC filtering on the FlexConnect AP or WLC and port security on the switch ensures that only the selected users access the wired network.

Most of the time, however, a company implements a WLAN to facilitate the mobility of the employees, which implies that a FlexConnect AP or WLC, at any given time, does not have a predetermined number of users associated with it.

Therefore, in cases where it is impossible to determine the number of users connected to the AP, enabling port security on the switch port offers no advantages. At worst, it can create an involuntary DoS attack if the policy for port security is set to shut down the port in the event of a violation. When this happens, all the users connected to that AP lose network connectivity. [Figure 4-30](#) shows an example of using port security to limit a wireless MAC flooding attack by locking down the port and sending an SNMP trap.

Figure 4-30 Using Port Security

Effectiveness of Port Security

When port security is not an option to stop an attack, a MAC flooding attack will not succeed if it is launched by a wireless user. The reason for this is the 802.11 protocol itself. Association with an AP is MAC-based; this means that the AP bridges (translational bridge) traffic coming from or going to known users (known MACs). If a MAC flooding attack is launched from a wireless user, all the 802.11 frames with random source MAC addresses that are not associated to the AP are dropped. The only frame allowed is the one with the MAC address of the malicious user, which the switch has probably already learned. Thus, the fundamental behavior of the access point itself prevents the switch from being susceptible to MAC flooding attacks.

Using Port Security to Mitigate a DHCP Starvation Attack

For wired access, port security can currently prevent a DHCP starvation attack launched from a PC connected to a switch that is using a tool such as Gobbler. The inability of the attack to succeed is due more to a limitation of the tool than the mitigation offered by port security. The only reason such an attack fails is that Gobbler uses a different source MAC address to generate a different DHCP request and can be mitigated by port protection.

However, if an attacker is able to use their MAC address in the Ethernet packet and simply changes the MAC address in the DHCP payload (the field is called chaddr), port security would not stop the attack. In this case, all that can currently be done is to try to slow down the attack using a DHCP rate limiter on the switch port.

Wireless DHCP Starvation Attack

In a Unified Wireless Network deployment, the vulnerability to a DHCP starvation attack differs between a WLC terminating the user traffic or a FlexConnect terminating the user traffic.

The WLC protects the network from DHCP starvation attacks because it examines DHCP requests to ensure that the client MAC address matches the chaddr. If the addresses do not match, the DHCP request is dropped.

In the case of FlexConnect, the user VLAN is terminated locally, the DHCP request does not go through the controller, and an analysis of the chaddr cannot be performed. In this case, the same security considerations apply for this method of access as they do for wired access. A smart (wireless) attacker uses the MAC address with which he or she is associated to the AP to generate the random DHCP requests, and then simply changes the requesting MAC address within the DHCP packet payload. To the AP, the packet looks valid because the originating MAC is the same as the MAC used to associate to the trusted AP.

DHCP Snooping to Mitigate a Rogue DHCP Server Attack

DHCP snooping is a DHCP security feature that provides security by building and maintaining a DHCP snooping binding table and filtering untrusted DHCP messages. It does this by differentiating between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCP server or another switch. End-user ports can be restricted to sending only DHCP requests and no other type of DHCP traffic. Trusted ports allow any DHCP message to be forwarded. The DHCP snooping table is built per VLAN and ties the IP address/MAC address of the client to the untrusted port. Enabling DHCP snooping prevents users from connecting a non-authorized DHCP server to an untrusted (user-facing) port and start replying to DHCP requests.

DHCP Snooping for Wireless Access


The WLC manages all DHCP requests from clients and acts as a DHCP relay agent. DHCP requests from WLAN clients are not broadcast back out to the WLAN, and they are unicasted from the WLC to a configured DHCP server. This protects other WLAN clients connected to the WLC from rogue DHCP server attacks.

Clients connecting to VLANs via an FlexConnect 802.1q trunk interface are not protected against rogue DHCP server attacks.

Keep in mind that the CISF features (in this case DHCP snooping) are implemented on the switch, not on the AP, so the ability of a switch to intercept malicious messages from a rogue server only happens if traffic is seen by the switch.

Figure 4-31 shows an example of using DHCP snooping to mitigate against a rogue DHCP server attack, and how an attack can happen before the switch is able to provide DHCP protection.

Figure 4-31 Security Used Against Rogue DHCP Server Attack



Effectiveness of DHCP Snooping

DHCP snooping is enabled on a per-VLAN basis, so it works on a trunk port. A separate DHCP snooping entry is inserted for each DHCP request received on a given trunk port for clients in different VLANs. The fact that DHCP snooping works on trunk ports is very important because it makes this CISF feature applicable to a WLAN deployment where multiple SSIDs/VLANs are configured on the local interface of the FlexConnect WLC. If an attacker is associated to the same WLAN/VLAN as the target, but via a different FlexConnect WLC, the switch is able to protect against the DHCP spoof attack. However, if the attacker and the target are associated to the same FlexConnect WLC, the attack does not traverse the access switch and it is not detected.

DHCP snooping also provides some protection against DHCP server attacks by rate limiting the DHCP requests to the DHCP server.

Dynamic ARP Inspection to Mitigate a Man-in-the-Middle Attack

Dynamic ARP Inspection (DAI) is enabled on the access switch on a per-VLAN basis. It compares ARP requests and responses, including gratuitous ARPs (GARPs), with the MAC/IP entries populated by DHCP snooping in the DHCP binding table. If the switch receives an ARP message with no matching entry in the DHCP binding table, the packet is discarded and a log message is sent to the console.

DAI prevents ARP poisoning attacks that can lead to man-in-the-middle (MIM) attacks such as those launched using ettercap. Ettercap stops the GARP messages sent by the malicious user to the target to alter the ARP table to receive the malicious user's traffic. The ARP messages are filtered directly at the port to which the attacker is connected.

DAI for Wireless Access

The WLC protects against MIM attacks by performing a similar function as DAI on the WLC itself. DAI should not be enabled on the access switch for those VLANs connecting directly to the WLCs because the WLC uses GARP to support Layer 3 client roaming.

It is possible to enable DAI for each VLAN configured on a trunk between a FlexConnect and access point. Therefore, DAI is useful in wireless deployments where multiple SSIDs/VLANs exist on an FlexConnect. However, in an FlexConnect WLC deployment, there are two topologies that impact the effectiveness of the DAI feature. Both topologies assume that the attacker is associated to a FlexConnect WLC and is Layer 2-adjacent to the targets:

- Topology 1—One target is wireless and associated to the same AP as the attacker, while the other target is the default gateway. This is the most common attack.
- Topology 2—Both targets are wireless.

The following examples illustrate how attacks are launched and stopped.

- The MIM attack attempts to use GARP to change the ARP table entries for the default gateway and the wireless target in order to redirect traffic to the attacker. DAI can block GARP for the default gateway, but DAI has no impact on a spoofed GARP for the wireless target. This limits the effectiveness of the MIM attack, but does not completely prevent the effects of the MIM attack.
- The MIM attack sends GARPs to wireless clients. The switch implementing DAI does not see these GARPs and cannot block the attack.

Figure 4-32 is an example of the attack mechanism where GARPs are sent to the two IP connection nodes on the subnet to divert the traffic between them.

Because the MAC address is provided in the log, the administrator can take further blocking action by disassociating the attacker.

When DAI is configured on a VLAN, an ARP rate limiter is configured globally to prevent flooding of ARP requests coming from a certain port. The default value of the rate limiter is 15 packets per second (pps). If this limit is reached, the switch disables the port to prevent the attack. In this case, to launch a MIM attack, an attacker must first discover who else is Layer 2 adjacent. To do this, ettercap generates a series of GARPs, claiming to be each one of the IP addresses on the subnet. In this way, the real owner of that address replies and ettercap can build its table.

In lab tests, this limit has been reached immediately when using ettercap and the port shuts down. This is acceptable in a wired topology. In a wireless topology, shutting down the port connected to the AP causes all the wireless users to lose their connection to the outside world and a possible MIM attack turns into a DoS attack.

To avoid this potential DoS (involuntarily created by enabling DAI), Cisco recommends turning off the ARP rate limiter on the port of the switch connected to the AP. You can do this with the following interface level command:

```
ip arp inspection limit none
```

An alternative is to change the threshold to a value larger than 15 pps. However, this is not a general remedy because it depends on the implementation of the specific tool being used to launch the attack.

Using IP Source Guard to Mitigate IP and MAC Spoofing

When enabled on an interface of the access switch, IP Source Guard dynamically creates a per-port access control list (PACL) based on the contents of the DHCP snooping binding table. This PACL enforces traffic to be sourced from the IP address issued at DHCP binding time and prevents any traffic from being forwarded by other spoofed addresses. This also prevents an attacker from impersonating a valid address by either manually changing the address or running a program designed to do address spoofing, such as hping2. This feature has an option (port security) to filter the incoming address, also using the MAC address in the DHCP snooping binding table.

The attacker typically uses the spoofed address to hide his or her real identity and launch an attack, such as a DoS attack, against a target.

IP Source Guard for Wireless Access

In the case of wireless access, IP Source Guard can be enabled on the trunk port connecting the access switch to the FlexConnect WLC. This allows the switch to filter any traffic coming from wireless users that does not match an entry in the DHCP binding table.

IP Source Guard does not need to be enabled on the VLANs configured behind a WLC, because the WLC performs a similar function to ensure that the IP address used by a client is the IP address that has been assigned to that client.


IP Source Guard is beneficial in FlexConnect WLC deployments because the FlexConnect AP (unlike a standard AP) is not able to check the WLAN client MAC-to-IP address binding relationship.

In tests, the following two topologies were considered:

- Topology 1—The target is represented by another wireless user associated to the same AP.
- Topology 2—The target is another wireless user associated to a different AP.

Figure 4-33 is an example of using IP Source Guard to mitigate IP and MAC spoofing attacks.

Figure 4-33 *IP Source Guard Preventing MIM*



Effectiveness of IP Source Guard

The effectiveness of the IP Source Guard feature depends on two factors: the way the attacker is able to spoof the address, and which topology is being tested.

An association to the AP is based on the client MAC address, so if the AP receives a frame with an unknown source MAC address, it drops the frame. When launching an IP spoofing attack, the attacker has the option to use his or her own MAC address or to use one from another user connected to the same AP. All the other combinations, such as using a random MAC address or using the MAC address of a user connected to another AP, lead to a failed attack because the AP drops the frame.

In case the attacker uses his or her own MAC address but spoofs the IP address, IP Source Guard enabled on the switch stops the attack in the second topology but not the first. In the first topology, the traffic stays local to the AP and the CISF feature is not invoked. In the other topologies, CISF successfully stops the attack because the IP-spoofed packet sent by the malicious user has no entry in the DHCP snooping table.

However, if the attacker is able to spoof both the MAC and the IP address of another wireless user connected to the same AP, basically assuming the identity of another user, the attack is successful in topologies 1 and 2. Spoofing both the MAC and IP address is realistically possible in a hotspot environment where no encryption is used, or when the weaknesses of WEP are exploited. This is one of the reasons why Cisco highly recommends the use of strong encryption whenever possible.

Summary of Target Attacks

Table 4-2 presents a summary of applicable target attacks, consideration and solutions.

Table 4-2 *Summary of Findings*

Targeted Attack	Applicability	Considerations	Solution
MAC flooding	No	Macof uses random MAC addresses as source and destination	AP discards frames from a source MAC not in the association table
Targeted Attack	Applicability	Considerations	Solution
DHCP starvation	Yes on FlexConnect Controller discards bad DHCP requests	The requesting MAC is carried in the DHCP payload	None—rate limiting
Rogue DHCP server	Yes on FlexConnect Controller blocks DHCP offers from the WLAN	It is assumed the rogue DHCP server is wireless	None
MIM between wireless clients	Yes on FlexConnect Controller blocks GARPs	Traffic does not go through the switch in this case	None
MIM between wireless clients on different APs	Yes on FlexConnect Controller blocks GARPs	The hacker can intercept traffic only toward the wire.	DAI with violation
MIM between wireless and wired clients	Yes on FlexConnect Not a supported controller configuration	The hacker can intercept traffic only toward the wire.	DAI with violation
IP spoofing	Yes on FlexConnect Controller checks IP address and MAC address binding	Encryption over the air is required to prevent identity spoofing	IP Source Guard



Note

Only those attacks that are targeted by the CISF features are on wired access, and it is always assumed that the attacker is wireless, while the target could be either wired or wireless depending on the topologies considered.



Cisco Unified Wireless QoS

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

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

QoS Overview

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

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

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

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

Wireless QoS Deployment Schemes

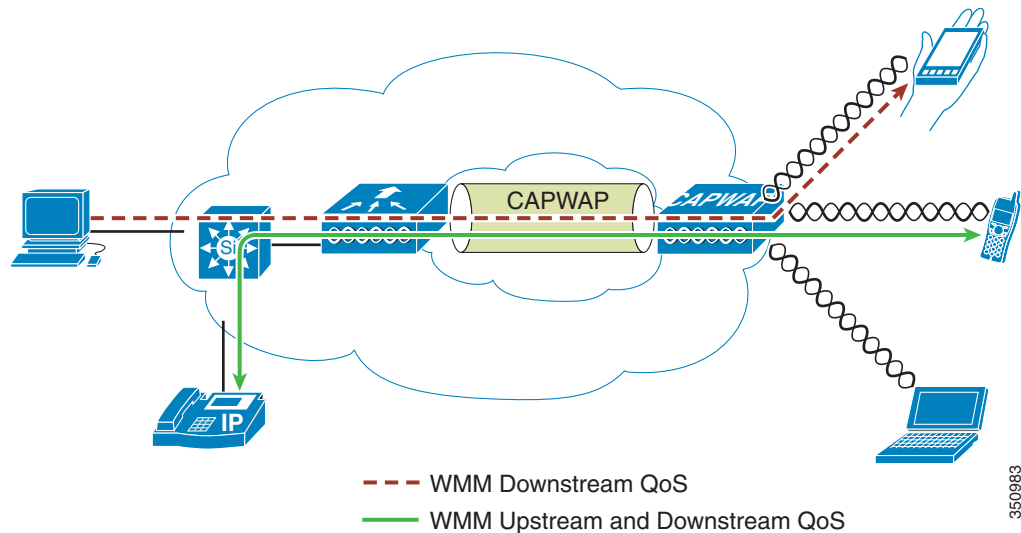
In the past, WLANs were mainly used to transport low-bandwidth, data-application traffic. Currently, with the expansion of WLANs into vertical (such as retail, finance, and education) and enterprise environments, WLANs are used to transport high-bandwidth data applications, in conjunction with time-sensitive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

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

Cisco Unified Wireless products support Wi-Fi MultiMedia (WMM), a QoS system based on IEEE 802.11e that has been published by the Wi-Fi Alliance and WMM Power Save, as well as Admission Control.

Figure 5-1 illustrates an example of the deployment of wireless QoS based on Cisco Unified Wireless technology features.

Figure 5-1 QoS Deployment Example



QoS Parameters

QoS is defined as the measure of performance for a transmission system that reflects its transmission quality and service availability. Service availability is a crucial component of QoS. Before QoS can be successfully implemented, the network infrastructure must be highly available.

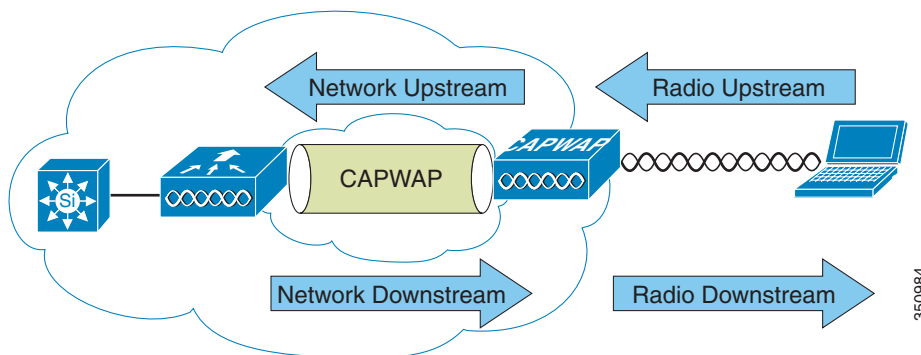
Network transmission quality is determined by the elements of latency, jitter, and loss, as shown in Table 5-1.

Table 5-1 QoS Transmission Quality

Element	Description
Latency	<p>Latency (or delay) is the amount of time it takes for a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time period is called the end-to-end delay and can be divided into two areas:</p> <ul style="list-style-type: none"> Fixed network delay—Includes encoding and decoding time (for audio and video), and the finite amount of time required for the electrical or optical pulses to traverse the media en route to their destination. Variable network delay—Generally refers to network conditions, such as queuing and congestion, that can affect the overall time required for transit.
Jitter	<p>Jitter (or delay-variance) is the difference in the end-to-end latency between packets. For example, if one packet requires 100 ms to traverse the network from the source endpoint to the destination endpoint, and the next packet requires 125 ms to make the same trip, the jitter is calculated as 25 ms.</p>
Loss	<p>Loss (or packet loss) is a comparative measure of packets successfully transmitted and received to the total number that were transmitted. Loss is expressed as the percentage of packets that were dropped.</p>

Radio Upstream and Downstream QoS

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

Figure 5-2 Upstream and Downstream QoS

As illustrated in Figure 5-2:

- *Radio downstream* QoS—Traffic leaving the AP and traveling to the WLAN clients. Radio downstream QoS is the primary focus of this chapter, because this is still the most common deployment. The radio client upstream QoS depends on the client implementation.
- *Radio upstream* QoS—Traffic leaving the WLAN clients and traveling to the AP. WMM provides upstream QoS for WLAN clients supporting WMM.

- *Network downstream*—Traffic leaving the wireless LAN controller (WLC) traveling to the AP. QoS can be applied at this point to prioritize and rate-limit traffic to the AP



Note Configuration of *Ethernet downstream* QoS is not described in this guide.

- *Network upstream*—Traffic leaving the AP, traveling to the WLC. The AP classifies traffic from the AP to the upstream network according to the traffic classification rules of the AP.

QoS and Network Performance

The application of QoS features could be difficult to detect on a lightly loaded network. If latency, jitter, and loss are noticeable when the media is lightly loaded, it indicates either a system fault, poor network design, or that the latency, jitter, and loss requirements of the application are not a good match for the network. QoS features start to be applied to application performance as the load on the network increases. QoS works to keep latency, jitter, and loss for selected traffic types within acceptable boundaries. When providing only radio downstream QoS from the AP, radio upstream client traffic is treated as best-effort. A client must compete with other clients for upstream transmission as well as competing with best-effort transmission from the AP. Under certain load conditions, a client can experience upstream congestion, and the performance of QoS-sensitive applications might be unacceptable despite the QoS features on the AP. Ideally, upstream and downstream QoS can be operated either by using WMM on both the AP and WLAN client, or by using WMM and a client proprietary implementation.



Note

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

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

802.11 Distributed Coordination Function

Data frames in 802.11 are sent using the distributed coordination function (DCF), which is composed of the following main components:

- Interframe spaces (IFS including SIFS, PIFS, and DIFS, which are described below)
- Random backoff (contention window)

DCF is used in 802.11 networks to manage access to the RF medium. A baseline understanding of DCF is necessary to deploy 802.11e-based enhanced distributed channel access (EDCA). For more information on DCF, see the IEEE 802.11 specification at:

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

These 802.11 DCF components are discussed further in the following sections.

Interframe Spaces

The 802.11 standard defines interframe spaces (IFS) as:

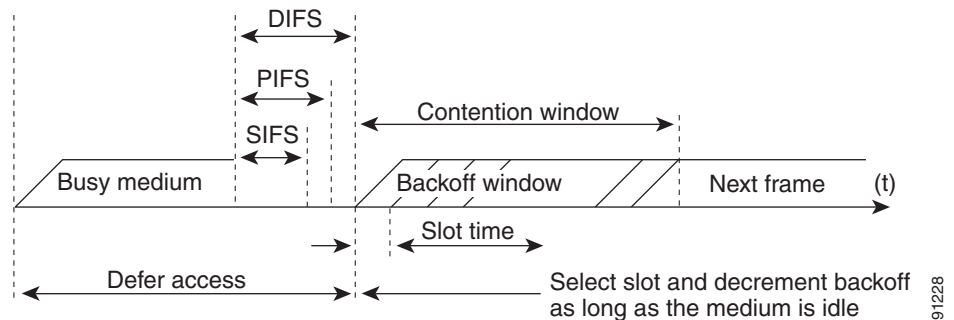
- Short interframe space (SIFS)—10 μ s
- PCF interframe space (PIFS)—SIFS + 1 x slot time = 30 μ s
- DCF interframe space (DIFS)—50 μ s SIFS + 2 x slot time = 50 μ s



Note The base timing used in the IFS example shown in Figure 5-3 is for 802.11b. The timing in 802.11g and 802.11a are different, but the principles applied are the same.

IFS allow 802.11 to control which traffic gets first access to the channel after carrier sense declares the channel to be free. Generally, 802.11 management frames and frames not expecting contention (a frame that is part of a sequence of frames) use SIFS, and data frames use DIFS, as shown in Figure 5-3.

Figure 5-3 Interframe Spaces




Random Backoff

When DCF has a data frame ready to be transmitted, the DCF goes through the following steps:

1. DCF generates a random backoff number between zero and a minimum contention window (see [aCWmin](#), [aCWmax](#), and [Retries](#), page 5-6).
2. DCF waits until the channel is free for a DIFS interval.
3. If the channel is still free, DCF begins to decrement the random backoff number for every slot time (20 μ s) that the channel remains free.
4. If the channel becomes busy (such as when a station gets to zero), DCF stops the decrement and steps 2 and 3 are repeated.
5. If the channel remains free until the random backoff number reaches zero, DCF allows the frame to be transmitted.

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

Figure 5-4 Distributed Coordination Function Example

91229

The DCF steps illustrated in [Figure 5-4](#) are:

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


The process continues as traffic arrives on the different stations.

aCWmin, aCWmax, and Retries

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

- aCWmin—Minimum contention window
- aCWmax—Maximum contention window

The random number used in the random backoff is initially a number between 0 and aCWmin. If the initial random backoff expires without successfully transmitting the frame, the station or AP increments the retry counter and doubles the value random backoff window size. This doubling in size continues until the size equals aCWmax. The retries continue until the maximum retries or time to live (TTL) is reached. This process of doubling the backoff window is often referred to as a *binary exponential backoff*, and is illustrated in [Figure 5-5](#) where the aCWmin is 2^5-1 , and increases to 2^6-1 , on the next backoff level, up to the aCWmax value of $2^{10}-1$.

Figure 5-5 Growth in Random Backoff Range with Retries**Note**

These values are for 802.11b implementations. Values can be different for different physical layer implementations.

Wi-Fi Multimedia

This section describes three important Wi-Fi multimedia (WMM) topics:

- WMM Access
- WMM Classification
- WMM Queues

WMM Access

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

WMM Classification

WMM uses the 802.1P classification scheme (part of the IEEE 802.1D MAC Bridges standard). This classification scheme has eight priorities that WMM maps to four access categories with WMM designations:

- AC_BK—Background
- AC_BE—Best effort
- AC_VI—Video
- AC_VO—Voice

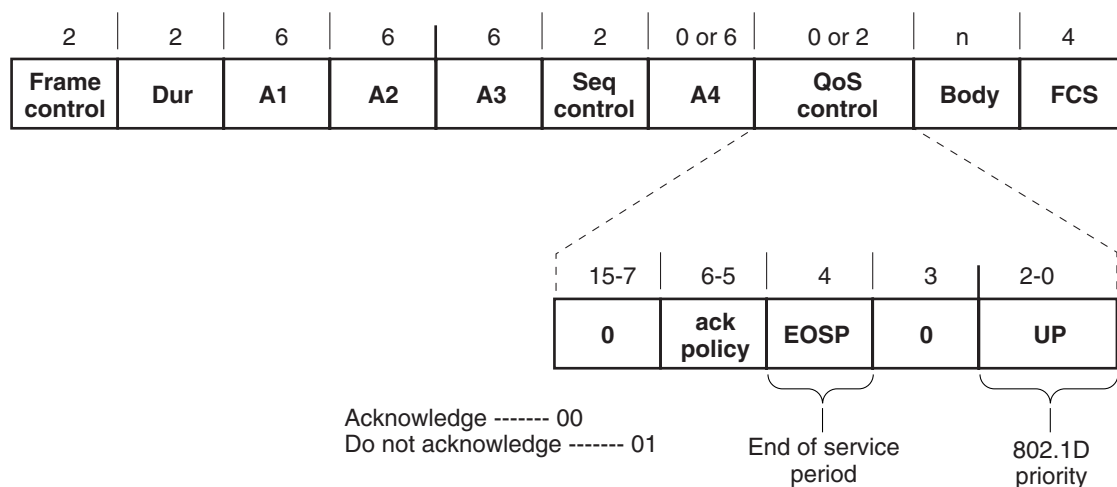
As shown in [Table 5-2](#), these access categories map to the four queues (see [WMM Queues, page 5-9](#)) required by WMM devices.

Table 5-2 **Table 2 802.1P and WMM Classification**

Priority	802.1P Priority	802.1P Designation	Access Category_WMM Designation
Lowest	1	BK	AC_BK
	2	-	
	0	BE	AC_BE
	3	EE	
	4	CL	AC_VI
	5	VI	
	6	VO	AC_VO
Highest	7	NC	

[Figure 5-6](#) shows the WMM data frame format. Note that even though WMM maps the eight 802.1P classifications to four access categories, the 802.1D classification is sent in the frame.

Figure 5-6 **WMM Frame Format**



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

WMM Queues

Figure 5-7 shows the queuing performed on a WMM client or AP. There are four separate queues, one for each of the access categories. Each of these queues contends for the wireless channel in a similar manner to the DCF mechanism described above, with each of the queues using different IFS, aCWmin, and aCWmax values. If more than one frame from different access categories collide internally, the frame with the higher priority is sent and the lower priority frame adjusts its backoff parameters as though it had collided with a frame external to the queuing mechanism. This system is called enhanced distributed channel access (EDCA).

Figure 5-7 WMM Queues

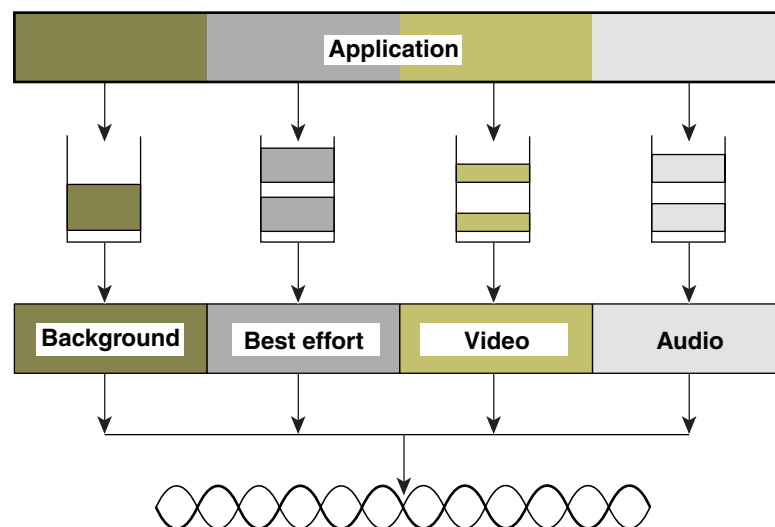



Figure 5-8 illustrates the principles behind EDCA, where different interframe spacing and aCWmin and aCWmax values (for clarity aCWmax is not shown) are applied per traffic classification. Different traffic types wait different IFS before counting down their random backoff. The aCW value used to generate the random backoff number also depends on the traffic classification. For example, the aCWmin[3] for Voice is 23-1, and aCWmin[5] for best-effort traffic is 25-1. High priority traffic has a small IFS and a small aCWmin value, giving a short random backoff, whereas best-effort traffic has a longer IFS and large aCWmin value that on average gives a large random backoff number.

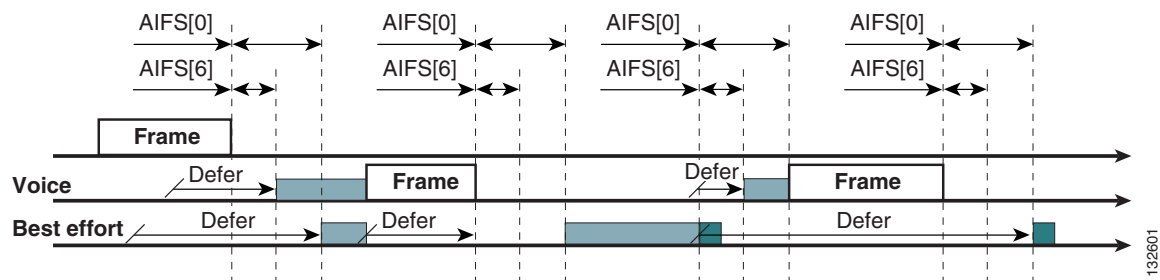
Figure 5-8 Access Category Timing



Enhanced Distributed Channel Access

Figure 5-9 illustrates an example of the enhanced distributed channel access (EDCA) process.

Figure 5-9 EDCA Example



The EDCA process follows the sequence:

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

The access category settings shown in [Table 5-3](#) and [Table 5-4](#) are, by default, the same for an 802.11a radio and are based on formulas defined in WMM.

**Note**

[Table 5-3](#) refers to the parameter settings on a client, which are slightly different from the settings for an AP. The AP has a larger AIFS[n] for audio and video admission controls (ACs).

Table 5-3 WMM Client Parameters

AC	CWmin	aCWmax	AIFS[n]	TXOP Limit (802.11b)	TXOP Limit (802.11a/g)
AC_BK	CWmin	aCWmax	7	0	0
AC_BE	CWmin	$4*(aCQmin+1)-1$	3	0	0
AC_VI	$(CWmin+1)/2-1$	CWmin	1	6.016 ms	3.008 ms
AC_VO	$(CWmin+1)/4-1$	$(CWmin+1)/2-1$	1	3.264 ms	1.504 ms

Table 5-4 WMM AP Parameters

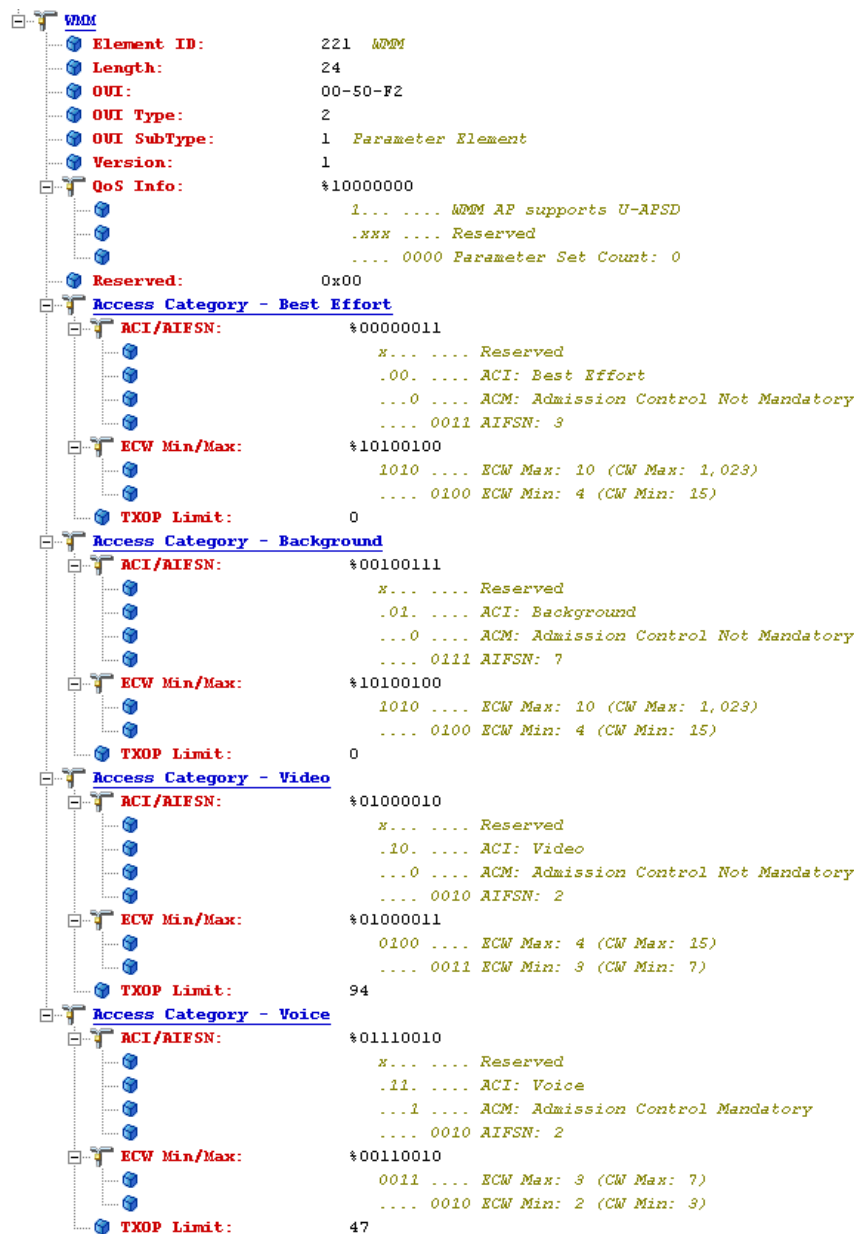
Access Category	CWmin	aCWmax	AIFS[n]	TXOP Limit (802.11b)	TXOP Limit (802.11a/g)
AC_BK	CWmin	aCWmax	7	0	0
AC_BE	CWmin	$4*(aCQmin+1)-1$	3	0	0
AC_VI	$(CWmin+1)/2-1$	CWmin	2	6.016 ms	3.008 ms
AC_VO	$(CWmin+1)/4-1$	$(CWmin+1)/2-1$	2	3.264 ms	1.504 ms

The overall impact of the different AIFS, CWmin, and aCWmax values is difficult to illustrate in timing diagrams because their impact is more statistical in nature. It is easier to compare the AIFS and the size of the random backoff windows, as shown in [Figure 5-8](#).

When comparing Voice and Background frames as examples, these traffic categories have CWmin values of 2^3-1 (7) and 2^5-1 (31), and AIFS of 2 and 7, respectively. This is an average delay of 5 $(2+7/1)$ slot times before transmitting an audio frame, and an average of 22 slot $(7+31/2)$ times for Background frame. Therefore, Voice frames are statistically much more likely to be sent before Background frames.

[Figure 5-10](#) shows the WMM information in a probe response. Apart from the WMM access-category information contained in this element, the client also learns which WMM categories require admission control. As can be seen in this example, the Voice admission control (AC) is set to mandatory. This requires the client to transmit the request to the AP, and have the request accepted, before it can use this AC. Admission control is further discussed in different parts of this chapter.

Figure 5-10 Probe Response WMM Element Information



22-1939

Unscheduled-Automatic Power-save Delivery

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

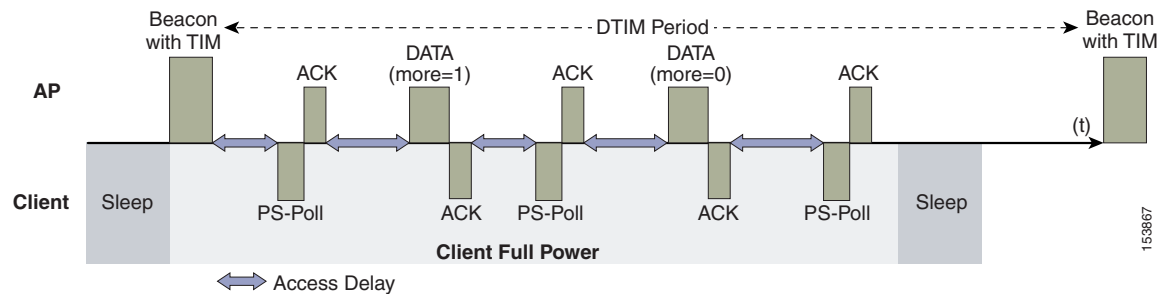
- The primary benefit of U-APSD is that it allows the audio client to synchronize the transmission and reception of audio frames with the AP, thereby allowing the client to go into power-save mode between the transmission/reception of each audio frame tuple. The WLAN client frame transmission in the access categories supporting U-APSD triggers the AP to transmit any data frames queued for that WLAN client in that access category. A U-APSD client continues listening to the AP until it receives a frame from the AP with an end-of-service period (EOSP) bit set. This tells the client that it can now go back into its power-save mode. This triggering mechanism is considered a more

efficient use of client power than the regular listening for beacons method, at a period controlled by the delivery traffic indication message (DTIM) interval. This is because the latency and jitter requirements of audio are such that a wireless VoIP client would either not be in power-save mode during a call, resulting in reduced talk times, or would use a short DTIM interval that results in reduced standby times. The use of U-APSD allows the use of long DTIM intervals to maximize standby time without sacrificing call quality. The U-APSD feature can be applied individually across access categories, allowing U-APSD can be applied to the audio ACs in the AP, but the other ACs still use the standard power-save mode feature.

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

Figure 5-11 shows a sample frame exchange for the standard 802.11 power save delivery process.


Figure 5-11 Standard Client Power-Save



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

This presents two major problems. The first is that it is quite inefficient, requiring the PS-polls, as well as the normal data exchange, to go through the standard access delays associated with DCF. The second issue, being more critical to audio traffic, is that retrieving the buffered data is dependent on the DTIM, which is a multiple of the beacon interval. Standard beacon intervals are 100 ms, and the DTIM interval can be integer multiples of this. This introduces a level of jitter that is generally unacceptable for audio calls, and audio handsets switch from power-save mode to full transmit and receive operation when a audio call is in progress. This gives acceptable audio quality but reduces battery life. The Cisco 7921G Unified Wireless IP Phone addresses this issue by providing a PS-Poll feature that allows the 7921G to generate PS-Poll requests without waiting for a beacon TIM. This allows the 7921G to poll for frames when it has sent a frame, and then go back to power-save mode. This feature does not provide the same efficiency as U-APSD, but improves battery life for 7921G phones on WLANs without U-APSD.

Figure 5-12 shows an example of traffic flows with U-APSD. In this case, the trigger for retrieving traffic is the client sending traffic to the AP. The AP, when acknowledging the frame, tells the client that data is queued for it and that it should stay connected. The AP then sends data to the client, typically as a TXOP burst where only the first frame has the EDCF access delay. All subsequent frames are then sent directly after the acknowledgment frame. In a VoWLAN implementation there is likely to be only one frame queued at the AP. The VoWLAN client is able to go into sleep mode after receiving that frame from the AP.

Figure 5-12 U-APSD

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

TSpec Admission Control

Traffic Specification (TSpec) allows an 802.11e client to signal its traffic requirements to the AP. In the 802.11e MAC definition, two mechanisms provide prioritized access: the contention-based EDCA option and the controlled access option provided by the transmit opportunity (TXOP). When describing TSpec features where a client can specify its traffic characteristics, it is easy to assume that this would automatically result in the use of the controlled access mechanism, and have the client granted a specific TXOP to match the TSpec request. However, this does not have to be the case; a TSpec request can be used to control the use of the various access categories (ACs) in EDCA. Before a client can send traffic of a certain priority type, it must have requested to do so by way of the TSpec mechanism. For example, a WLAN client device wanting to use the audio access categories must first make a request for use of that AC. Whether or not AC use is controlled by TSpec requests is configurable with audio and audio ACs controlled by TSpec requests, and best-effort and background ACs can be open for use without a TSpec request. The use of EDCA ACs, rather than the 802.11e Hybrid Coordinated Channel Access (HCCA), to meet TSpec requests is possible in many cases because the traffic parameters are sufficiently simple to allow them to be met by allocating capacity, rather than creating a specific TXOP to meet the application requirements.

Add Traffic Stream

The Add Traffic Stream (ADDTS) function is used by WLAN client to send an *admission request* to an AP. Signaling its TSpec request to the AP, an admission request is in one of two forms:

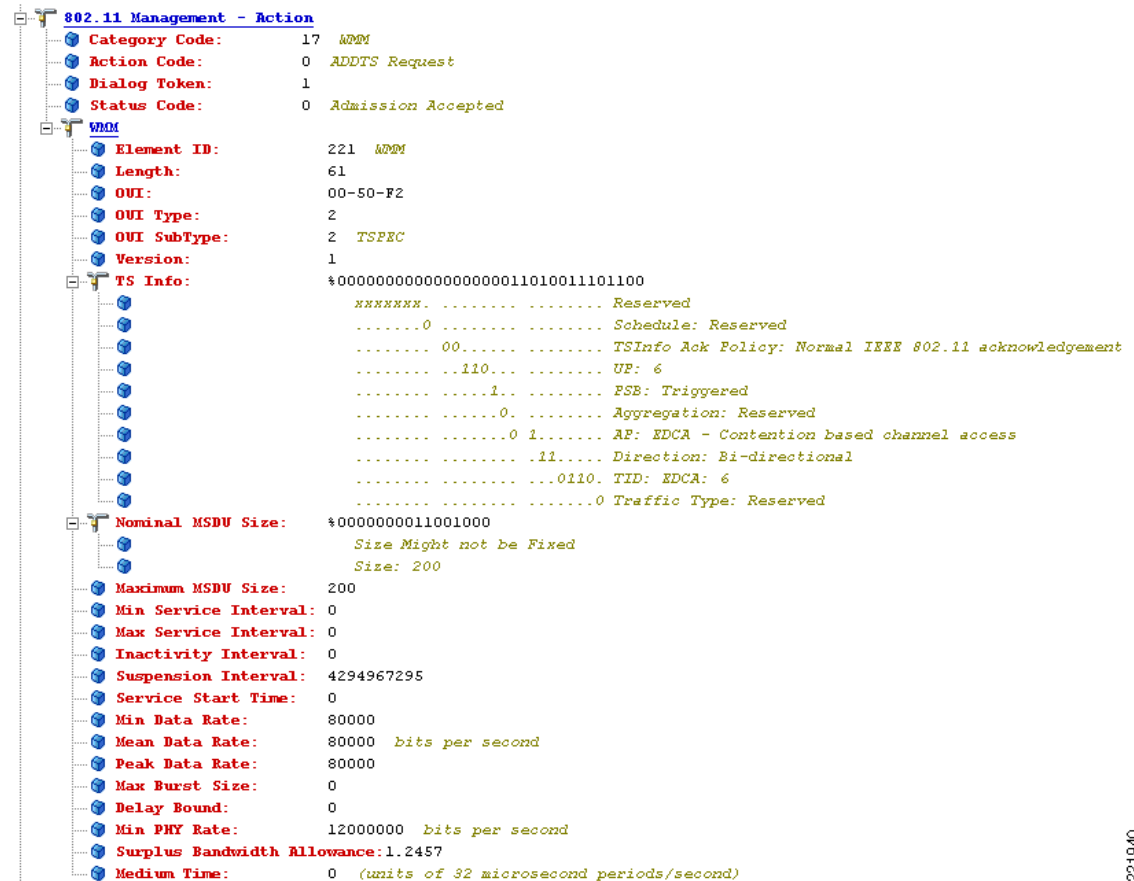
- ADDTS action frame—Created when a phone call is originated or terminated by a client associated to the AP. The ADDTS contains TSpec and could contain a traffic stream rate set (TSRS) information element (IE).
- Association and re-association message—The association message might contain one or more TSpecs and one TSRS IE if the station wants to establish the traffic stream as part of the association. The re-association message might contain one or more TSpecs and one TSRS IE if a station roams to another AP.

The ADDTS contains the TSpec element that describes the traffic request. See [Figure 5-13](#) and [Figure 5-14](#) for examples of an ADDTS request and response between a Cisco 7921 WLAN handset and a Cisco AP. Apart from key data describing the traffic requirements, such as data rates and frame sizes, the TSpec element also tells the AP the minimum physical rate that the client device will use. This allows the calculation of how much time that station can potentially consume in transmitting and receiving in this TSpec, and therefore allowing the AP to calculate whether it has the resources to meet the TSpec.

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

TSpec support is not required by clients. But when a WLAN is configured with call admission control (CAC) for either audio or video that client that is not in support of TSpec is must send the audio and video packets at a Best effort QoS level (see [QoS Profiles](#), page 5-16). So, if the WLAN is set at QoS level of audio or video and CAC is enabled then the correct behavior for a client without ADDTS logic is to send the audio and video traffic with Best effort markings. If a TSpec capable clients has its ADDTS request reject be the Wi-Fi channel utilization is high than the configured CAC limit. That client per specification is supposed to mark the audio and video packets at Best effort.

Figure 5-13 ADDTS Request Decode



221940

```

802.11 Management - Action
  Category Code: 17 WPM
  Action Code: 1 ADDTS Response
  Dialog Token: 1
  Status Code: 0 Admission Accepted
  WMM
    Element ID: 221 WPM
    Length: 61
    OUI: 00-50-F2
    OUI Type: 2
    OUI SubType: 2 TSPEC
    Version: 1
    TS Info: %00000000000000000000000011010011101100
      xxxxxxxx. .... Reserved
      .....0 ..... Schedule: Reserved
      ..... 00..... TSInfo Ack Policy: Normal IEEE 802.11 acknowledgement
      ..... .110..... UP: 6
      ..... .1..... PSE: Triggered
      ..... .0..... Aggregation: Reserved
      ..... .0 1..... AP: EDCA - Contention based channel access
      ..... .11..... Direction: Bi-directional
      ..... .0110. TID: EDCA: 6
      ..... .0 Traffic Type: Reserved
    Nominal MSDU Size: %0000000011001000
      Size Might not be Fixed
      Size: 200
    Maximum MSDU Size: 200
    Min Service Interval: 0
    Max Service Interval: 0
    Inactivity Interval: 0
    Suspension Interval: 4294967295
    Service Start Time: 0
    Min Data Rate: 80000
    Mean Data Rate: 80000 bits per second
    Peak Data Rate: 80000
    Max Burst Size: 0
    Delay Bound: 0
    Min PHY Rate: 12000000 bits per second
    Surplus Bandwidth Allowance: 1.2457
    Medium Time: 528 (units of 32 microsecond periods/second)

```

In addition to the WMM support described above, the Cisco *Centralized WLAN Architecture* has a number of advanced QoS features. These features include:

- QoS Profiles
- WMM Policy
- Voice over IP Phones
- Admission Control Parameters

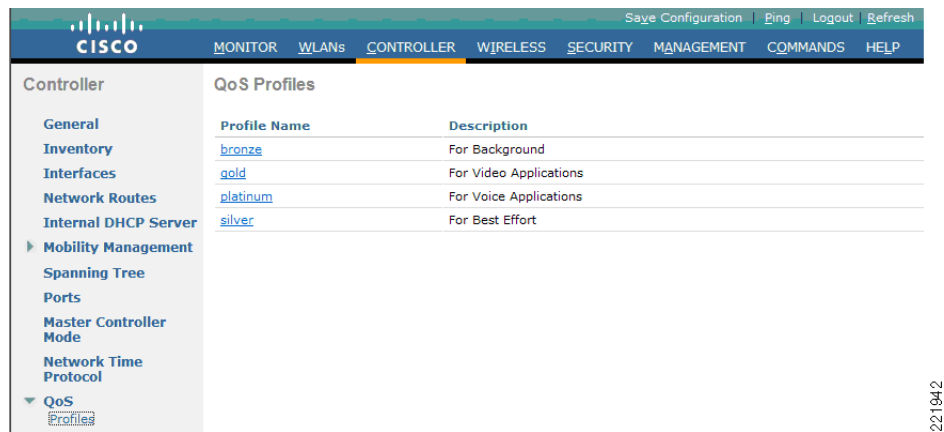
QoS Profiles

Primary among these are the QoS profiles used by the WLC. As shown in [Figure 5-15](#), the QoS profiles can be configured as:

- Bronze—Background
- Gold—Video applications

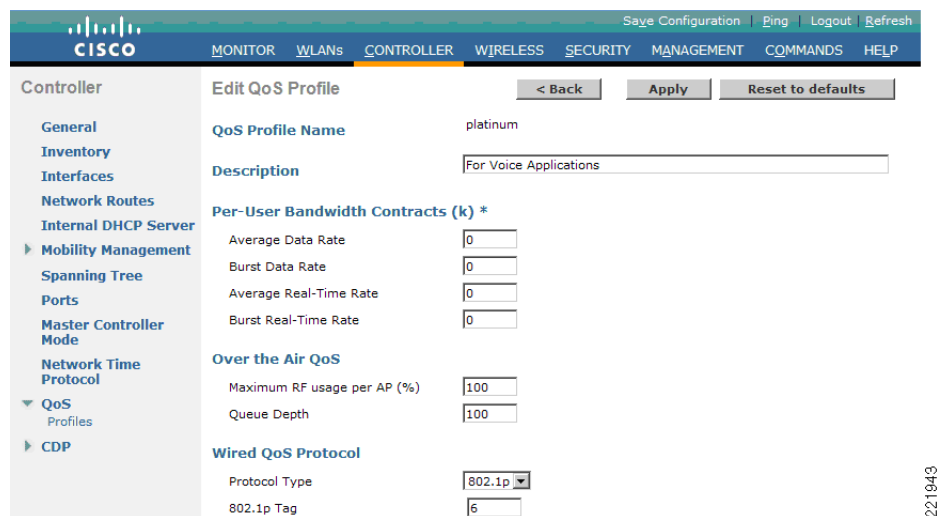
- Platinum—Voice applications
- Silver—Best effort

Figure 5-15 QoS Profile Options



Each of the profiles shown in Figure 5-16 allows the configuration of bandwidth contracts, RF usage control, and the maximum 802.1P classification allowed.

Figure 5-16 QoS Profile Settings



Cisco generally recommends that the Per-User Bandwidth Contracts settings be left at their default values and that the 802.11 WMM features be used to provide differentiated services.

For WLANs using a given profile, the 802.1P classification in that profile controls two important class of service (CoS) behaviors:

- Determines what CoS value packets initiated from the WLC are marked with.

The value of the CoS parameter is used to mark the CoS of all CAPWAP (*Control And Provisioning of Wireless Access Points*) packets for the WLAN using that profile. So a WLAN with a platinum QoS profile, and the 802.1P mark of 6, will have its CAPWAP packets from the application manager interface of the controller marked with CoS of 5. The WLC adjusts the CoS to be compliant with Cisco QoS baseline recommendations. The reason why it is important to maintain the IEEE CoS

marking in the configuration is below. If the WLAN is configured to trust CoS rather than DSCP at the network connection to the WLC, the CoS value is used for the DSCP of the CAPWAP packets received by the AP; and eventually the WMM classification and queuing for WLAN traffic. This is because the WLAN WMM classification of a frame is derived from the DSCP value of the CAPWAP packet carrying that frame.

- Determines the maximum CoS value that can be used by clients connected to that WLAN.

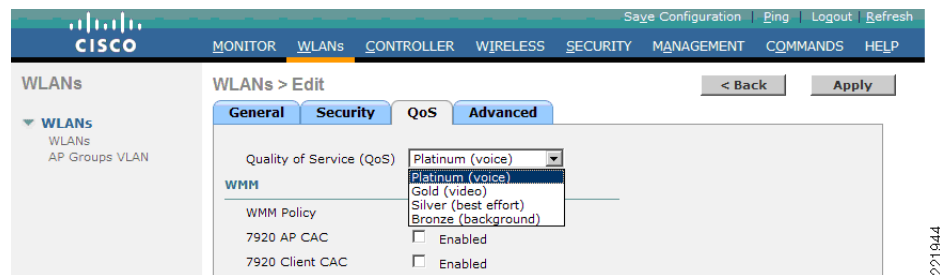
The 802.1P classification sets the maximum CoS value that is admitted on a WLAN with that profile.

WMM audio traffic arrives with a CoS of 6 at the AP, and the AP automatically performs a CoS-to-DSCP mapping for this traffic based on a CoS of 6. If the CoS value in the WLC configuration is set to a value less than 6, this changed value is used by the WLAN QoS profile at the AP to set the maximum CoS marking used and therefore which WMM admission control (AC) to use.

The key point is that with the Cisco Unified Wireless Network, you should always think in terms of IEEE 802.11e classifications and allow the Unified Wireless Network Solution to take responsibility for converting between IEEE classification and the Cisco QoS baseline.

The WLAN can be configured with various default QoS profiles, as shown in [Figure 5-17](#). Each of the QoS profiles are annotated with their typical use. In addition, clients can be assigned a QoS profile based on their identity, through authentication, authorization and accounting (AAA). For a typical enterprise, WLAN deployment parameters such as per-user bandwidth contracts and over-the-air QoS, should be left at their default values, and standard QoS mechanisms, such as WMM and wired QoS, should be used to provide optimum QoS to clients.

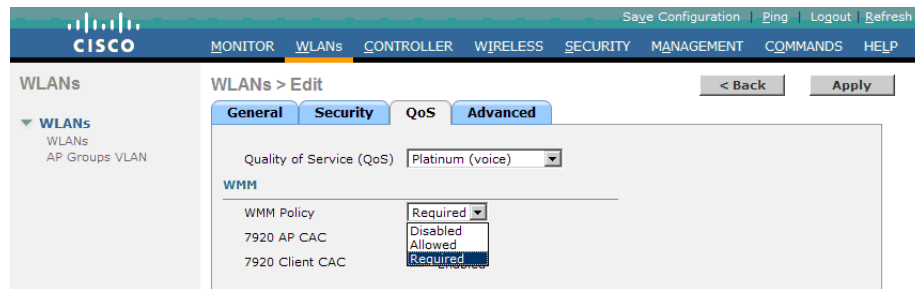
Figure 5-17 WLAN QoS Profile



WMM Policy

In addition to QoS profiles, WMM Policy for the WLAN allows you to control additional WMM options, as shown in [Figure 5-18](#). The WMM options are:

- Disabled—The WLAN does not advertise WMM capabilities nor allow WMM negotiations
- Allowed—The WLAN does allow WMM and non-WMM clients
- Required—Only WMM-enabled clients can be associated with this WLAN

Figure 5-18 WLAN WMM Policy

221945

Voice over IP Phones

Figure 5-19 shows the basic QoS Enhanced Basis Service Set (QBSS) information element (IE) advertised by a Cisco AP. The Load field indicates the portion of available bandwidth currently used to transport data on that AP.

Figure 5-19 QBSS Information Element

1 Octet	1 Octet	4 bytes
Element ID (11)	Length	Load

153873

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

- Old QBSS—Draft 6 (pre-standard)
- New QBSS—Draft 13 802.11e (standard)
- New distributed CAC load IE—A Cisco information element

The QBSS used depends on the WMM and Cisco 792x VoIP phone settings on the WLAN.

792x phone support, as shown in Figure 5-20, is a component of the WLC WLAN configuration that enables the AP to include the appropriate QBSS element in its beacons. WLAN clients with QoS requirements, such as Cisco 792x phones, use these advertised QoS parameters to determine the best AP with which to associate.

The WLC provides 792x phone support through the client call admission control (CAC) limit. This support includes:

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

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

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

- If 7920 AP CAC limit is to be supported, the number 3 QBSS IE is sent in the beacons and probe responses for bg radios.

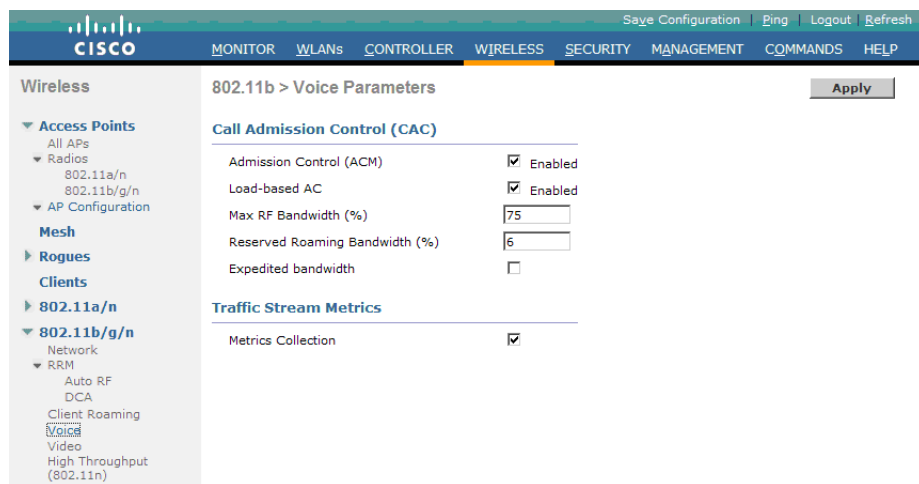
**Note**

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

Admission Control Parameters

Figure 5-20 shows an example of the configuration window for setting the Voice parameters on the controller.

Figure 5-20 Voice Parameter Setting



The CAC parameters include the *Max RF Bandwidth (%)* that a radio can be using and still accept the initiation of a VoWLAN call through a normal ADDTS request. The range of that value is 5 to 85 percent of the channel bandwidth.

The *Reserved Roaming Bandwidth (%)* parameter specifies how much capacity is reserved to be able to respond to ADDTS requests during association or re-association, and which are VoWLAN clients with calls in progress that are trying to roam to that AP.

To enable AC based upon these parameters, select the *Admission Control (ACM)* check box. This enables AC based upon the capacity of the AP but it does not take into account the possible *channel loading* impact of other APs in the area. To include this channel loading in capacity calculations, select the both *Load-Based AC* and *Admission Control (ACM)* check boxes.

**Note**


Voice and video load-based CAC applies to non-mesh APs. For mesh APs, only static CAC is applicable.

SIP CAC support requires either static or load-based CAC. If you are using *Static* CAC then SIP CAC support allows the configuration of the number of calls on the AP. Generally the dynamic the load-balanced approach is the better way of managing quantity of calls to prevent the quality from suffering from over subscription of calls on the Wi-Fi channel.

In the Voice Parameters window (Figure 5-20), the *Metrics Collection* option specifies whether data is collected on audio or video calls for use by Cisco Prime Infrastructure.

Figure 5-21 shows an example of one of the audio statistics reports available with Cisco Prime Infrastructure. The example shows the calls established on the radio of one AP and the number of calls that roamed to that AP. This report and other audio statistics can be scheduled or performed on request (ad-hoc) and either displayed graphically in Cisco Prime Infrastructure or written to a file.

Figure 5-21 Voice Statistics from Cisco Prime Infrastructure



Note

CAC is performed only for audio and video QoS profiles.

Figure 5-21 shows the effect of having a low percent of bandwidth set aside for audio CAC calls. Only enough bandwidth was reserved for four calls, but the calls were able to roam to other Wi-Fi channels. Figure 5-22 shows CAC options for media streaming. *Max RF Bandwidth* is shared between the audio, video and media streaming. The Voice, Video, and Media tabs each have their own *Max RF Bandwidth* that are added together for an aggregated total of the complete bandwidth reservation for media on a Wi-Fi channel. While each tab shows a maximum value of 85 percent for the field, the overall Max RF Bandwidth value is actually the sum of all three fields. If Max RF Bandwidth in the Voice tab is set to 85 percent then in video tab and media tabs the Max RF Bandwidth fields must be set to zero percent. If you wanted some bandwidth with CAC behavior on audio, video and data, then you could set the value to 25 percent in the fields of each tab. This would have a channel bandwidth limit for media of 75 percent. With each media type getting one quarter of the bandwidth and with data getting one fourth (1/4) of the bandwidth.

Figure 5-22 WLC 802.11a(5 GHz) Media Window

CAC for video behaves like audio CAC. The purpose of CAC for video is to limit the amount of video calling so that the quality of active video calls is not negatively impacted by additional video being added to the Wi-Fi channel.

**Note**

See the WLC configuration guide for more details on these and the other configuration options.

Impact of TSpec Admission Control

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

Table 5-5 and Table 5-6 describe the impact on classification if admission control is enabled or not and whether or not a traffic stream has been established.

Table 5-5 Upstream Traffic

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

Table 5-6 Downstream Traffic

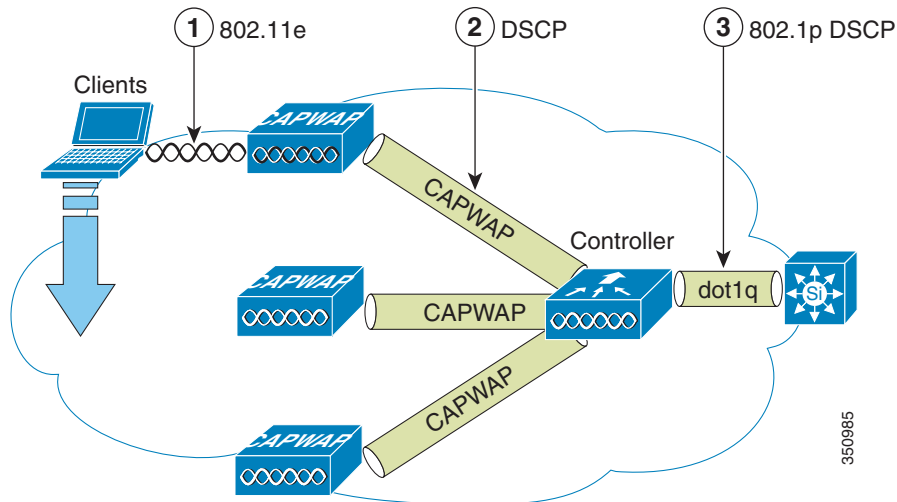
AC Enabled	Traffic Stream Established	No Traffic Stream
No	No change	No change
Yes	No change	Remark UP to BE for WMM client. For non-WMM clients, use WLAN QoS.

802.11e, 802.1P and DSCP Mapping

WLAN data in a Unified Wireless Network is tunneled by way of CAPWAP (IP UDP packets). To maintain the QoS classification that has been applied to WLAN frames, the WLC uses a process of mapping classifications to and from DSCP and CoS. For example, when WMM classified traffic is sent by a WLAN client, it has an 802.1P classification in its frame. The AP needs to translate this classification into a DSCP value for the CAPWAP packet carrying the frame to ensure that the packet is treated with the appropriate priority on its way to the WLC. A similar process must occur on the WLC for CAPWAP packets going to the AP.

A mechanism to classify traffic from non-WMM clients is also required so that their CAPWAP packets can also be given an appropriate DSCP classification (see [Classification Considerations, page 5-29](#)) by the AP and the WLC.

[Figure 5-23](#) shows the various classification mechanisms in the CAPWAP WLAN network.

Figure 5-23 WMM and 802.1P Relationship

Multiple classification mechanisms and client capabilities require multiple strategies. These strategies include:

- CAPWAP control frames require prioritization so they are marked with a DSCP classification of CS6 (an IP routing class).
- WMM-enabled clients have the classification of their frames mapped to a corresponding DSCP classification for CAPWAP packets to the WLC. This mapping follows the standard IEEE CoS-to-DSCP mapping, with the exception of the changes necessary for QoS baseline compliance. This DSCP value is translated at the WLC to a CoS value on 802.1Q frames leaving the WLC interfaces.
- Non-WMM clients have the DSCP of their CAPWAP tunnel set to match the default QoS profile for that WLAN. For example, the QoS profile for a WLAN supporting 792x phones would be set to platinum, resulting in a DSCP classification of EF for data frames packets from that AP WLAN.
- CAPWAP data packets from the WLC have a DSCP classification that is determined by the DSCP of the wired data packets sent to the WLC. The 802.11e classification used when transmitting frames from the AP to a WMM client is determined by the AP table converting DSCP to WMM classifications.

**Note**

The WMM classification used for traffic from the AP to the WLAN client is based on the DSCP value of the CAPWAP packet, and not the DSCP value of the contained IP packet. Therefore, it is critical that an end-to-end QoS system be in place.

QoS Baseline Priority Mapping

The CAPWAP AP and WLC perform QoS baseline conversion so that WMM values, as described in [Table 5-7](#), are mapped to the appropriate QoS baseline DSCP values, rather than the IEEE values.

Table 5-7 Access Point QoS Translation Values¹

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

1. The IEEE 802.11e UP (user priority) value for DSCP values that are not mentioned in the table is calculated by considering 3 MSB bits of DSCP. For example, the IEEE 802.11e UP value for DSCP 32 (100 000 in binary), would be the decimal converted value of the MSB (100) which is 4. The 802.11e UP value of DSCP 32 is 4.

Deploying QoS Features on CAPWAP-based APs

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

- The wired CAPWAP AP interface reads or writes Layer 2 CoS (802.1P) information. The WLC and the APs depend on Layer 3 classification (DSCP) information to communicate WLAN client traffic classification. This DSCP value could be subject to modification by intermediate routers, and therefore the Layer 2 classification received by the destination might not reflect the Layer 2 classification marked by the source of the CAPWAP traffic.
- The APs no longer use NULL VLAN ID. As a consequence, Layer 2 CAPWAP does not effectively support QoS because the AP does not send the 802.1P/Q tags and in Layer 2 CAPWAP there is no outer DSCP on which to fall back.
- APs do not re-classify frames; they prioritize them based on CoS value or WLAN profile.
- APs carry out EDCF-like queuing on the radio egress port only.
- APs do FIFO queuing only on the Ethernet egress port.

WAN QoS and FlexConnect

For WLANs that have data traffic forwarded to the WLC, the behavior is same as non-hybrid remote edge FlexConnect APs. For locally-switched WLANs with WMM traffic, FlexConnect APs mark the dot1p value in the dot1q VLAN tag for upstream traffic. This occurs only on tagged non-native VLANs.

For downstream traffic, FlexConnect APs use the incoming dot1q tag from the Ethernet side and then use this to queue and mark the WMM values on the radio of the locally-switched VLAN.

The WLAN QoS profile is applied to both upstream and downstream packets. For downstream traffic, if an 802.1P value that is higher than the default WLAN value is received, the default WLAN value is used. For upstream traffic, if the client sends an WMM value that is higher than the default WLAN value, the default WLAN value is used. For non-WMM traffic there is no CoS marking on the client frames from the AP.

Guidelines for Deploying Wireless QoS

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

QoS LAN Switch Configuration Example

AP Switch Configuration

The QoS configuration of the AP switch is minor because the switch must trust the DSCP of the CAPWAP packets that are passed to it from the AP. There is no CoS marking on the CAPWAP frames coming from the AP. Below is an example of this configuration. Note that this configuration addresses only the classification and that queuing commands can be added depending on local QoS policy.

```
interface GigabitEthernet1/0/1
  switchport access vlan 100
  switchport mode access
  mls qos trust dscp
  spanning-tree portfast
end
```

In trusting the AP DSCP values, the access switch is trusting the policy set for that AP by the WLC. The maximum DSCP value assigned to client traffic is based on the QoS policy applied to the WLAN on that AP.

WLC Switch Configuration

The QoS classification decision at the WLC-connected switch is slightly more complicated than at the AP-connected switch because the choice can be to either trust the DSCP or the CoS of traffic coming from the WLC. When making this decision, consider the following:

- Traffic leaving the WLC can be either upstream (to the WLC or network) or downstream (to the AP and WLAN client). The downstream traffic is CAPWAP encapsulated, and the upstream traffic is either CAPWAP encapsulated or decapsulated WLAN client traffic leaving the WLC.
- DSCP values of CAPWAP packets are controlled by the QoS policies on the WLC; the DSCP values set on the WLAN client traffic (encapsulated by the CAPWAP tunnel header) has not been altered from those set by the WLAN client.
- CoS values of frames leaving the WLC are set by the WLC QoS policies, regardless of whether they are upstream, downstream, encapsulated, or decapsulated.

The following example chooses to trust the CoS settings of the WLC because this allows a central location for the management of WLAN QoS rather than having to manage the WLC configuration and an additional policy at the WLC switch connection.

```
interface GigabitEthernet1/0/13
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 11-13,60,61
 switchport mode trunk
 mls qos trust cos
end
```

If you want to have a more precise degree of control you can implement QoS classification policies on the WLAN-client VLANs.

Traffic Shaping, Over the Air QoS, and WMM Clients

Traffic shaping and over-the-air QoS are useful tools in the absence of WLAN WMM features, but they do not address the prioritization of 802.11 traffic directly. For WLANs that support WMM clients or 792x handsets, the WLAN QoS mechanisms of these clients should be relied on; no traffic shaping or over-the-air QoS should be applied to these WLANs.

WLAN Voice and Cisco Phones

The data sheets for Cisco Unified Communication Endpoints can be found at:

http://www.cisco.com/en/US/prod/voicesw/ps6788/ip_phones.html


For a general overview of Cisco Jabber, see:

<http://www.cisco.com/web/products/voice/jabber.html>

CAPWAP over WAN Connections

This section describes QoS strategies when CAPWAP APs are deployed across WAN links, as shown in Figure 5-24.

Figure 5-24 CAPWAP Traffic Across the WAN



CAPWAP Traffic Classification

CAPWAP APs can be generally separated into the following two types:

- CAPWAP control traffic—Identified by UDP port 5246
- CAPWAP 802.11 traffic—Identified by UDP port 5247

CAPWAP Control Traffic

CAPWAP control traffic can be generally divided into the following two additional types:

- Initialization traffic—Generated when a CAPWAP AP is booted and joins a CAPWAP system. For example, initialization traffic could be generated by controller discovery, AP configuration, and AP firmware updates.

**Note**

CAPWAP image packets from the controller are marked best effort, but their acknowledgement is marked CS6. Note that no sliding window protocol is used and each additional packet is sent only after an acknowledgement. This type of handshaking minimizes the impact of downloading files over a WLAN.

- Background traffic—Generated by an CAPWAP AP when it is an operating member of a WLAN network. Examples included CAPWAP heartbeat, radio resource management (RRM), and rogue AP measurements. Background CAPWAP control traffic is marked CS6.

Figure 5-24 show an example of an initial CAPWAP control message. The list of initial CAPWAP control messages includes:

- CAPWAP discovery messages
- CAPWAP join messages
- CAPWAP configuration messages
- Initial CAPWAP RRM messages

Figure 5-25 CAPWAP Discovery Request on a WISM-2

```

0 Frame 1: 162 bytes on wire (1304 bits), 162 bytes captured (1295 bits)
on Ethernet II, Src: Cisco_3a:ff:61 (04:7d:4f:3a:ff:61), Dst: broadcast (ff:ff:ff:ff:ff:ff)
Internet Protocol, Src: 10.30.0.130 (10.30.0.130), Dst: 255.255.255.255 (255.255.255.255)
  Version: 4
  Header length: 20 bytes
  Differentiated Services Field: 0xc0 (DSCP 0x10: Class Selector 6; DCN: 0x00)
  Total Length: 148
  Identification: 0x0000 (0)
  Flags: 0x02 (Don't Fragment)
  Fragment offset: 0
  Time to live: 255
  Protocol: UDP (17)
  Header checksum: 0x0598 [correct]
  Source: 10.30.0.130 (10.30.0.130)
  Destination: 255.255.255.255 (255.255.255.255)
  User Datagram Protocol, Src Port: 45048 (45048), Dst Port: capwap-control (5246)
    Source port: 45048 (45048)
    Destination port: capwap-control (5246)
    Length: 128
    Checksum: 0x0000 (none)
  Control And Provisioning of Wireless Access Points
    Preamble
      Version: 0
      Type: CAPWAP Header (0)
    Header
      Header Length: 4
      Radio ID: 0
      Wireless Binding ID: 16EE 802.11 (1)
    Header Flags
      Fragment ID: 0
      Fragment offset: 0
      Reserved: 0
      MAC length: 6
      MAC address: Cisco_49:fe:40 (04:fe:7f:49:fe:40)
      Padding for 4 Byte Alignment: 4b
    Control header

```

CAPWAP 802.11 Traffic

CAPWAP 802.11 traffic can be divided generally into the following two additional types:

- 802.11 management frames—802.11 management frames such as probe requests and association requests/responses are classified automatically with a DSCP of CS6.
- 801.11 data frames—Client data and 802.1X data from the client is classified according to the WLAN QoS settings, but packets containing 802.1X frames from the WLC are marked CS4. 802.11 data traffic classification depends on the QoS policies applied in the WLAN configuration and is not automatic. The default classification for WLAN data traffic is Best effort.

Classification Considerations

The DSCP classification used for CAPWAP control traffic is CS6 (an IP routing class) and is intended for IP routing protocols such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and others.

The current CAPWAP DSCP classification represents a classification that, although optimal for the WLAN system, might not align with your QoS policies and needs.

In particular, you might want to minimize the amount of CS6-classified traffic generated by the WLAN network. You might want to stop CS6 traffic generated by client activity such as probe requests. The easiest way to do this is to reclassify the CAPWAP 802.11 CS6 traffic to be a DSCP value with a lower QoS priority. The fact that the CAPWAP UDP port used is different from that used by CAPWAP data, and the default DSCP marking, allow for remarking this traffic without resorting to deep packet inspection.

In addition, you might want to ensure that CAPWAP initialization traffic does not impact routing traffic. The easiest way to ensure this is to mark with a lower priority the CAPWAP control traffic that is in excess of the background rate.

Router Configuration Examples

This section provides examples of router configurations that you can use as guides when addressing CS6 remarking or CAPWAP control traffic load.

The examples use CAPWAP APs on the 192.168.101.0/24 subnet and two WLCs with AP managers at 192.168.60.11 and 192.168.62.11.

Remarking Client Generated CS6 Packets

The following example shows a router configuration for remarking CAPWAP data packets marked as CS6 to a more appropriate value of CS3. This moves the traffic to a more suitable classification, at the level of call control, rather than at the level of network control.

```
class-map match-all CAPWAPDATA6
  match access-group 110
  match dscp cs6
!
policy-map CAPWAPDATA6
  class CAPWAPDATA6
    set dscp cs3
!
interface FastEthernet0
  ip address 192.168.203.1 255.255.255.252
  service-policy input CAPWAPDATA6
!
access-list 110 remark CAPWAP Data
access-list 110 permit udp 192.168.101.0 0.0.0.255 host 192.168.60.11 eq 5247
access-list 110 permit udp 192.168.101.0 0.0.0.255 host 192.168.62.11 eq 5247
access-list 111 remark CAPWAP Control
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.60.11 eq 5246
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.62.11 eq 5246
```

Changing the DSCP of CAPWAP Control Traffic above a predefined rate

The following is an example of rate limiting the CAPWAP control traffic from the WAN site to minimize the impact of the CS6-marked control traffic on routing traffic. Note that the rate limit configuration does not drop non-conforming traffic, but simply reclassifies that traffic.



Note

The following is an example and not a recommendation. Under normal circumstances, and following the design guidelines for deploying APs over a WAN connection, it is unlikely that CAPWAP control traffic would impact the WAN routing protocol connection.

```
interface Serial0
  ip address 192.168.202.2 255.255.255.252
  rate-limit output access-group 111 8000 3000 6000 conform-action transmit exceed-action
  set-dscp-transmit 26
access-list 111 remark CAPWAP Control
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.60.11 eq 5246
access-list 111 permit udp 192.168.101.0 0.0.0.255 host 192.168.62.11 eq 5246
!
```

For more information on WLAN QoS and 802.11e, see the *IEEE 802.11 Handbook: A Designer's Companion, 2nd Edition*, by Bob O'Hara and Al Petrick. ISBN: 978-0-7381-4449-8



Cisco Unified Wireless Multicast Design

Introduction

This chapter describes the Cisco Unified Wireless Multicast in IP multicast forwarding and provides information on how to deploy multicast in a wireless environment. A prerequisite for using the multicast performance functionality is that a multicast-enabled network must be configured on all routers between the controllers and the Access Points (APs). To accommodate networks that do not support multicast, the controller continues to support the original unicast packet forwarding mechanism.

IP multicast is a delivery protocol for information to a group of destinations. It uses the most efficient strategy to deliver the information over each link of the network. It sends only one copy of the information at each hop of the network, creating copies only when the links to the destinations split. Typically, many of today's networks applications use unicast packets i.e., from one source to one destination. However, when multiple receivers require the same data, replicating the data from the source to all the receivers as individual unicast packets increases the network load. IP multicast enables efficient transfer of data from a set of sources to a dynamically formed set of receivers.

IP multicast is typically used today for one way streaming media, such as video to large groups of receivers. Many cable TV operators, educational institutions and large enterprises have deployed IP multicast for their content delivery needs. Additionally, there have been some uses of audio and video conferencing using multicast. Another widespread use of multicast within campus and commercial networks is for file distribution, particularly to deliver operating system images and updates to remote hosts. IP multicast has also seen deployment within the financial sector for applications such as stock tickers and hoot-n-holler systems.

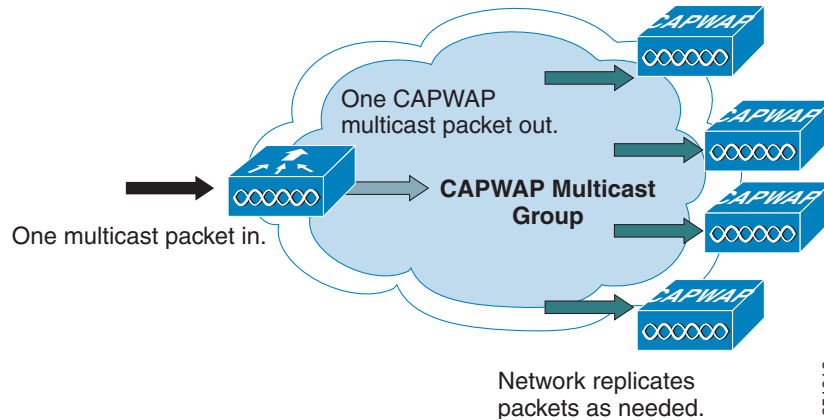
Overview of Multicast Forwarding

With Cisco Unified Wireless Network Software Release 4.1, significant enhancements were made to support the effective use of multicast in a wireless network. In 3.1 and prior software releases, packets intended to be multicast were actually unicast on the wireless network. Multicast support was added in 3.2, but there were some configuration limitations that required broadcast to be enabled. With 4.1 controller software release, separate broadcast and multicast support was enabled, allowing networks to be configured with just multicast, broadcast, or the use of both multicast and broadcast.

With the current Cisco Unified Wireless multicast support, each multicast frame received by the controller from a VLAN on the first hop router is copied and sent to the multicast group configured on the controller for the AP that is associated, as shown in [Figure 6-1](#). The multicast CAPWAP packet containing the multicast packet uses a WLAN bitmap, which tells the receiving AP which WLAN it must

forward the packet to. When the AP receives the CAPWAP packet, it strips off the outer CAPWAP encapsulation and transmits the multicast packet to the WLAN (on all radios associated to the WLAN) identified in the CAPWAP WLAN ID bitmask.

Figure 6-1 Multicast Forwarding Mechanism in Version 4.1 and Below



Effectively, a CAPWAP multicast group is used to deliver the multicast packet to each access point. This allows the routers in the network to use standard multicast techniques to replicate and deliver multicast packets to the APs. For the CAPWAP multicast group, the controller becomes the multicast source and the APs become the multicast receivers.



Note

A prerequisite for using the multicast performance functionality is that a multicast enabled network is configured on all routers between the controllers and the APs. To accommodate networks that do not support multicast, the controller continues to support the original unicast packet forwarding mechanism.



Note

With multicast enabled, any kind of multicast packet received on the VLAN from the first hop router is transmitted over the wireless including HSRP hellos, all router, EIGRP, and PIM multicast packets.

After the administrator enables multicast (multicast mode is disabled by default) and configures a CAPWAP multicast group, the access points' download the controller's CAPWAP multicast group address during the normal join process (at boot time) to the controller. After an access point joins a controller and downloads its configuration, the AP issues an Internet Group Management Protocol (IGMP) join request to join the controller's CAPWAP multicast group. This triggers the normal setup for the multicast state in the multicast-enabled routers between the controller and APs. The source IP address for the multicast group is the controller's management interface IP address, not the AP-manager IP address used for Layer 3 mode. Once the AP has joined the controllers CAPWAP multicast group, the multicast algorithm for client multicast traffic works as described below.

When the source of the multicast group is on the wired LAN:

- When the controller receives a multicast packet from any of the client VLANs on the first hop router, it transmits the packet to the CAPWAP multicast group via the management interface at the best effort QoS classification. The QoS bits for the CAPWAP multicast packet are hard coded at the lowest level and are not user changeable.
- The multicast-enabled network delivers the CAPWAP multicast packet to each of the access points that have joined the CAPWAP multicast group, using the normal multicast mechanisms in the routers to replicate the packet along the way as needed so that the multicast packet reaches all APs (Figure 6-1). This relieves the controller from replicating the multicast packets.

- Access points may receive other multicast packets but will only process the multicast packets that are sourced from the controller they are currently joined to; any other copies are discarded. If more than one WLAN is associated to the VLAN interface where the original multicast packet was sourced, the AP transmits the multicast packet over each WLAN (following the WLAN bitmap in the CAPWAP header). Additionally, if that WLAN is on both radios (802.11g and 802.11a), both radios transmit the multicast packet on the WLAN if there are clients associated, even if those clients did not request the multicast traffic.

When the source of the multicast group is a wireless client:

- The multicast packet is unicast (CAPWAP encapsulated) from the AP to the controller similar to standard wireless client traffic.
- The controller makes two copies of the multicast packet. One copy is sent out the VLAN associated with the WLAN it came on, enabling receivers on the wired LAN to receive the multicast stream and the router to learn about the new multicast group. The second copy of the packet is CAPWAP-encapsulated and is sent to the CAPWAP multicast group so that wireless clients may receive the multicast stream.


Wireless Multicast Roaming

A major challenge for a multicast client in a wireless environment is maintaining its multicast group membership when moving about the WLAN. Drops in the wireless connection moving from AP-to-AP can cause a disruption in a client's multicast application. Internet Group Management Protocol (IGMP) plays an important role in the maintenance of dynamic group membership information.

A basic comprehension of IGMP is important for understanding what happens to a client's multicast session when it roams about the network. In a Layer 2 roaming case, sessions are maintained simply because the foreign AP, if configured properly, already belongs to the multicast group and traffic is not tunneled to a different anchor point on the network. Layer 3 roaming environments are a little more complex in this manner and depending on what tunneling mode you have configured on your controllers, the IGMP messages sent from a wireless client will be affected. The default mobility tunneling mode on a controller is asymmetrical. As discussed in the [Chapter 2, "Cisco Unified Wireless Technology and Architecture,"](#) this means that return traffic to the client is sent to the anchor WLC then forwarded to the foreign WLC where the associated client connection resides. Outbound packets are forwarded out the foreign WLC interface. In symmetrical mobility tunneling mode, both inbound and outbound traffic are tunneled to the anchor controller. For more information on mobility tunneling, see [Chapter 2, "Cisco Unified Wireless Technology and Architecture."](#)

Asymmetric Multicast Tunneling

In asymmetric multicast tunneling, when a client roams to a new AP associated to a different WLC and on a different subnet, it is queried for its multicast group memberships by the foreign WLC and send out an IGMP group membership report. This is forwarded out the foreign WLC dynamic interface assigned to the VLAN and the client rejoins the multicast stream through the foreign subnet. [Figure 6-2](#) illustrates the traffic flow for normal data and multicast data.

Figure 6-2 Asymmetric Tunneling**Note**

In the event of a client roam, there is a slight disruption in the multicast session; in some applications it might be considered unsuitable for use.

Multicast Enabled Networks

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

CAPWAP Multicast Reserved Ports and Addresses

The controller blocks all multicast packets sent to any multicast group that have a destination port of 5246, 5247, and 5248. Additionally, all packets with a multicast group address equal to the controller's CAPWAP multicast group address are blocked at the controller. This is to prevent fragmented CAPWAP encapsulated packets from another controller being retransmitted (see the [Fragmentation and CAPWAP Multicast Packets](#) for more information). Ensure that the multicast applications on your network do not use these reserved ports or CAPWAP multicast group addresses.

Enabling Multicast Forwarding on the Controller

IP Multicast traffic through the controller is disabled by default. WLAN clients cannot receive multicast traffic when it is disabled. If you wish to turn on multicast traffic to the WLAN clients, follow these steps:

- Step 1** If you *have* a multicast enabled network, select **multicast** under Ethernet Multicast Mode to use the method where the network replicates the packets
- Step 2** You *do not have* a multicast enabled network, select **unicast** under Ethernet Multicast Mode to use the method where the controller replicates the packets.
- Step 3** On the controller general webpage, ensure the CAPWAP transport mode is set to Layer 3. The multicast performance feature only works in this mode.
- Step 4** Select multicast in the drop down menu for the Ethernet Multicast Mode and type in a multicast group address. This option is shown in [Figure 6-3](#).

Figure 6-3 Commands to turn on Ethernet Multicast Mode via the GUI.

The screenshot shows the Cisco Unified Wireless Multicast GUI. The top navigation bar includes links for MONITOR, WLANS, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, and HELP. The left sidebar lists various configuration sections under the Controller tab, with 'Mobility Management' selected. The main content area is titled 'General' and contains several configuration options:

- 802.3x Flow Control Mode: Disabled
- LWAPP Transport Mode: Layer 3 (Current Operating Mode is Layer3)
- LAG Mode on next reboot: Disabled (LAG Mode is currently disabled).
- Ethernet Multicast Mode: Multicast (239.255.1.57 is entered in the adjacent field, highlighted with a red box. Below it, the text 'Multicast Group Address' and 'H-REAP supports 'unicast' mode only.' are visible.)
- Broadcast Forwarding: Disabled
- Aggressive Load Balancing: Disabled
- Peer to Peer Blocking Mode: Disabled
- Over The Air Provisioning of AP: Enabled
- AP Fallback: Enabled
- Apple Talk Bridging: Disabled

An 'Apply' button is located at the top right of the configuration area. The page number 222290 is visible in the bottom right corner.

CLI Commands to Enable Ethernet Multicast Mode

- Step 1** Enable the CLI command: **configure network multicast global enable**
- Step 2** Enable the CLI command: **config network multicast mode multicast <IP Address>**
 Use the **show network** command to verify the multicast mode on the controller and **show capwap mcast** to verify the group on the AP. Other useful commands are **show ip mroute** and **show ip igmp membership** on the routers.

Multicast Deployment Considerations

Recommendations for Choosing a CAPWAP Multicast Address

**Caution**

Although not recommended, any multicast address can be assigned to the CAPWAP multicast group including the reserved link local multicast addresses used by OSPF, EIGRP, PIM, HSRP, and other multicast protocols.

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

**Note**

Do not use the 239.0.0.X address range or the 239.128.0.X address range. Addresses in these ranges overlap with the link local MAC addresses and will flood out all switch ports even with IGMP snooping turned on.

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

For more information on multicast address guidelines, refer to the document at the following URL:

http://www.cisco.com/en/US/tech/tk828/technologies_white_paper09186a00802d4643.shtml

Fragmentation and CAPWAP Multicast Packets

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

There are two different options to prevent this from happening, either of which is effective by itself. One, you may assign all controllers to the same CAPWAP multicast group address. Or two, you can apply standard multicast filtering techniques to ensure that CAPWAP encapsulated multicast packets do not reach any other controller. All Controllers have the Same CAPWAP Multicast Group lists the pros and cons of these two techniques.

Figure 6-4 *Pros and Cons of using the same Multicast Group or Different Groups*

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

All Controllers have the Same CAPWAP Multicast Group

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

However, assigning every controller to the same CAPWAP multicast group creates other problems. IGMP version 1 and 2 used by the APs to join the CAPWAP multicast group use Any Source Multicast (ASM) and the APs will receive multicast traffic from all sources of the multicast group in the network. This means the APs will receive multicast packets from all of the controllers on the network if the controllers are configured with the same multicast group address, and no multicast boundaries have been applied. One controller's multicast traffic will flood out to all of the APs across the network and every APs receive (and drop it if the source address is not equal to its controller's management address) the multicast traffic that is being received from any wireless multicast client in the entire network. Additionally, locally sourced multicast packets from any client VLAN such as HSRP, PIM, and EIGRP and OSPF multicast packets will also be flooded throughout the network.



Note

Cisco IOS APs (for example, 1240) use IGMPv2 while VxWorks APs (for example, 1030) use IGMPv1.

Controlling Multicast on the WLAN Using Standard Multicast Techniques

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



Note

A wired client anywhere in the network may request the CAPWAP multicast stream and receive it from all sources (if multicast boundaries are not applied). Multicast streams are not encrypted when they are encapsulated in the CAPWAP multicast packet. Therefore, it is recommended that multicast boundaries be implemented to block this type of access.

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

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

A typical example on a wireless client VLAN is given below. For more information on other multicast commands for a multicast enabled network refer to <http://www.cisco.com/go/multicast>. Filtering for multicast-enabled traffic also allows you to prevent propagation of certain worms like the sasser worm which relied on the TCP and ICMP transports with multicast addresses. Blocking these types of traffic with multicast group addresses does not affect most applications since they typically use UDP or TCP for streaming.

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

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

How Controller Placement Impacts Multicast Traffic and Roaming



Note

The multicast stream in either deployment, distributed or collocated, is not rate-limited and there is no way to put ACLs on it. Once enabled, all multicast traffic is forwarded to the wireless LAN including HSRP, EIGRP, OSPF, and PIM packets.

We look at two different deployments (distributed and centralized) and how they impact roaming with multicast clients. In a centralized deployment, WLC WLAN interfaces are attached to the same VLANs/subnets, the multicast stream is uninterrupted when a multicast client roams from APs on one WLC to an AP on another WLC. The centralized deployment creates a flat WLC client multicast network. The reason centralized WLCs do not affect multicast roaming is because once the multicast stream is requested from a single multicast client on a WLAN, it streams out all APs on that WLAN, on all radios (802.11g and 802.11a), on all WLCs, even if that access point WLAN has no clients associated with it that have requested the multicast traffic. If you have more than one WLAN associated to the VLAN, the AP transmits the multicast packet over each WLAN. Both the unicast mode CAPWAP packet and the multicast mode CAPWAP packet contain a WLAN bitmap that tells the receiving AP which WLAN it must forward the packet over.

The distributed deployment does not have this problem because while the WLANs are the same, the WLCs are attached to different VLANs. This means that when the multicast client roams to a new WLC, the WLC will first query the client for its multicast group memberships. At this point the client responds with its group membership report and the WLC forwards this message to the appropriate multicast group address through the VLAN associated with its local VLAN. This allows the client to resume its multicast session through the foreign WLC.

The distributed deployment reduces the amount of multicast traffic on the APs because, although the WLAN SSIDs are the same, the WLCs are attached to different VLANs. WLAN multicast traffic depends on a client request on the VLAN of that WLC. [Table 6-1](#) lists the advantages and disadvantages of distributed and collocated deployments.

Table 6-1 *Pros and Cons of Centralized WLCs and Distributed WLCs*

Deployment	PRO	CON
All centralized WLC WLANs connected to the same VLANs (subnets)	Multicast traffic started on any client VLAN will be transmitted to all APs so clients roaming to any AP will receive multicast stream	If only one client requests multicast traffic, all APs attached to all controllers will receive the stream and transmit it if they have any clients associated even if those clients did not request the multicast stream
Distributed WLCs on different VLANs and subnet	Multicast streams are isolated to APs attached to controller	Disruptions caused by multicast stream establishments after client roam

Additional Considerations

Two areas for additional consideration in multicast deployment are when implementing AP groups, and FlexConnect and APs. AP groups allow APs on the same controller to map the same WLAN (SSID) to different VLANs. If a client is roaming between APs in different groups, the multicast session will not function properly as this is currently not supported. Currently, the WLC forwards multicast only for the VLAN configured on the WLAN and does not take into consideration VLANs configured in AP groups.

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

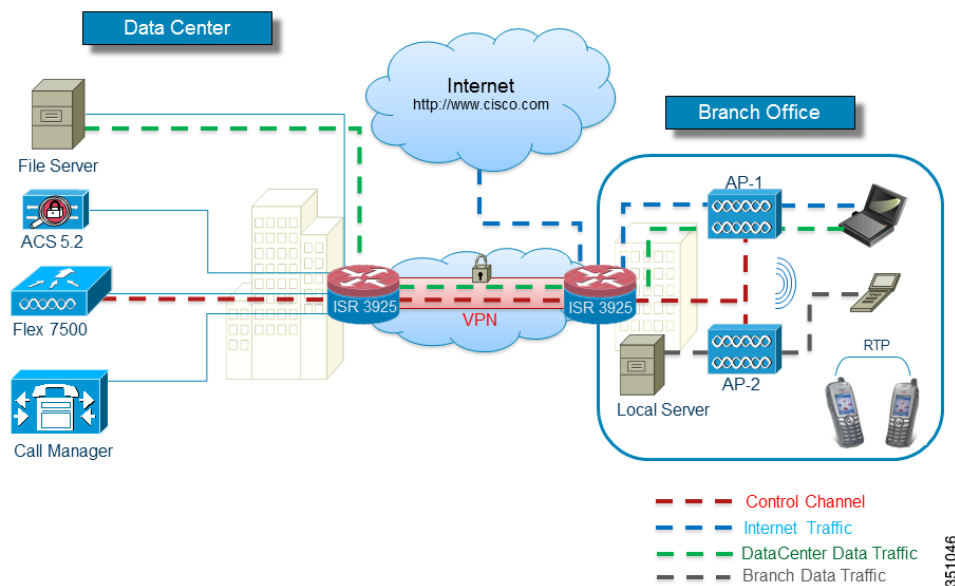
Even if the CAPWAP multicast packets are not able to transit the network to the FlexConnect AP, WLAN clients on that FlexConnect AP are able to send IGMP joins to the network connected to the WLC, as these are unicast messages.



FlexConnect

FlexConnect (previously known as Hybrid Remote Edge Access Point or H-REAP) is a wireless solution for branch office and remote office deployments. It enables you to configure and control access points in a branch or remote office from the corporate office through a wide area network (WAN) link without the deployment of a controller in each office. The FlexConnect access points (APs) can switch client data traffic locally and perform client authentication locally. When they are connected to the controller, they can also send traffic back to the controller.

Figure 7-1 FlexConnect Architecture



Note

To view the FlexConnect feature matrix, see:

http://www.cisco.com/en/US/products/ps10315/products_tech_note09186a0080b3690b.shtml#matrix

Supported Platforms

FlexConnect is only supported on these components:

- 1130AG, 1140, 1240AG, 1040, 1250, 1260, 1600, 2600, 3600, AP801, 3500I, 3500E and AP 1260 access points
- Cisco Flex 7500, Cisco 8500, 5500, 4400, and 2500 Series Controllers
- Catalyst 3750G Integrated Wireless LAN Controller Switch
- Cisco WiSM-2
- Controller Network Module for Integrated Services Routers
- Cisco virtual controller

FlexConnect Terminology

For clarity, this section provides a summary of the FlexConnect terminology and definitions used throughout this chapter.

Switching Modes

FlexConnect APs are capable of supporting the following switching modes concurrently, on a per-WLAN basis.

Local Switched

Locally-switched WLANs map wireless user traffic to discrete VLANs via 802.1Q trunking, either to an adjacent router or switch. If so desired, one or more WLANs can be mapped to the same local 802.1Q VLAN.

A branch user, who is associated to a local switched WLAN, has their traffic forwarded by the on-site router. Traffic destined off-site (to the central site) is forwarded as standard IP packets by the branch router. All AP control/management-related traffic is sent to the centralized Wireless LAN Controller (WLC) separately via Control and Provisioning of Wireless Access Points protocol (CAPWAP).

Central Switched

Central switched WLANs tunnel both the wireless user traffic and all control traffic via CAPWAP to the centralized WLC where the user traffic is mapped to a dynamic interface/VLAN on the WLC. This is the normal CAPWAP mode of operation.

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

Operation Modes

There are two modes of operation for the FlexConnect AP.

Connected mode—The WLC is reachable. In this mode the FlexConnect AP has CAPWAP connectivity with its WLC.

Standalone mode—The WLC is unreachable. The FlexConnect has lost or failed to establish CAPWAP connectivity with its WLC: for example, when there is a WAN link outage between a branch and its central site.

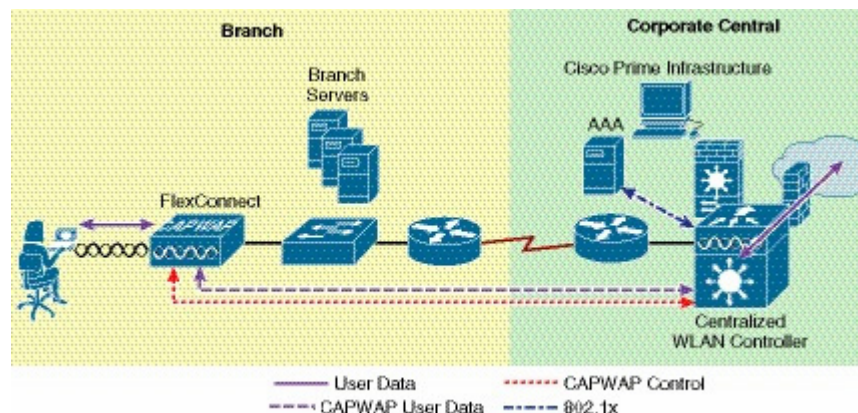
FlexConnect States

A FlexConnect WLAN, depending on its configuration and network connectivity, is classified as being in one of the following defined states.

Authentication-Central/Switch-Central

This state represents a WLAN that uses a centralized authentication method such as 802.1X, VPN, or web. User traffic is sent to the WLC via CAPWAP. This state is supported only when FlexConnect is in connected mode (Figure 7-2); 802.1X is used in the example, but other mechanisms are equally applicable.

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



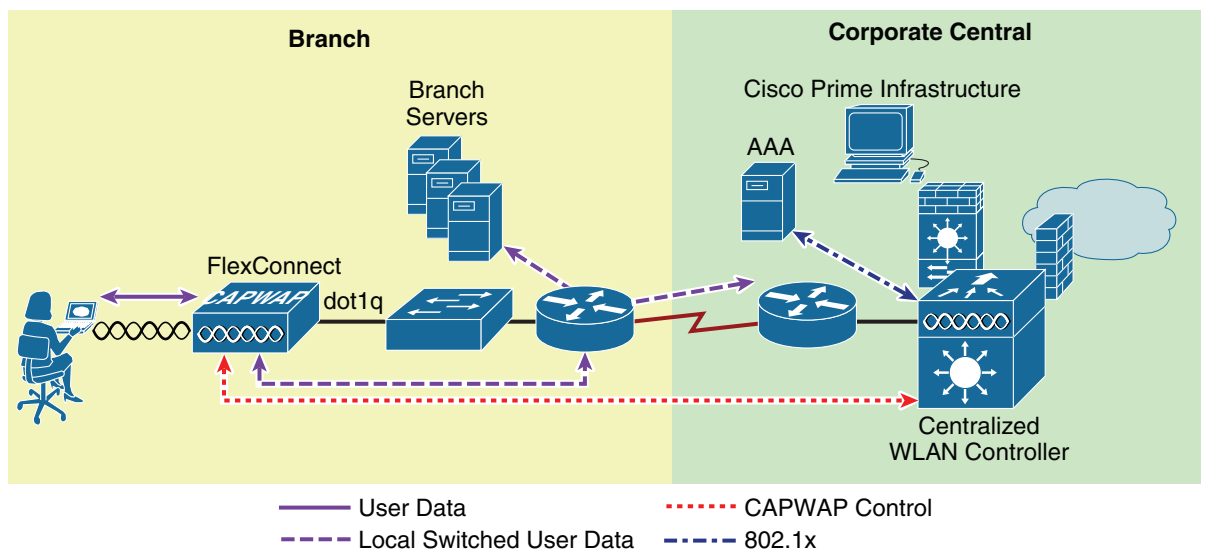
Authentication Down/Switching Down

Central switched WLANs (above) no longer beacon or respond to probe requests when the FlexConnect AP is in standalone mode. Existing clients are disassociated.

Authentication-Central/Switch-Local

This state represents a WLAN that uses centralized authentication, but user traffic is switched locally. This state is supported only when the FlexConnect AP is in connected mode (Figure 7-3); 802.1X is used in the Figure 7-3 example, but other mechanisms are equally applicable.

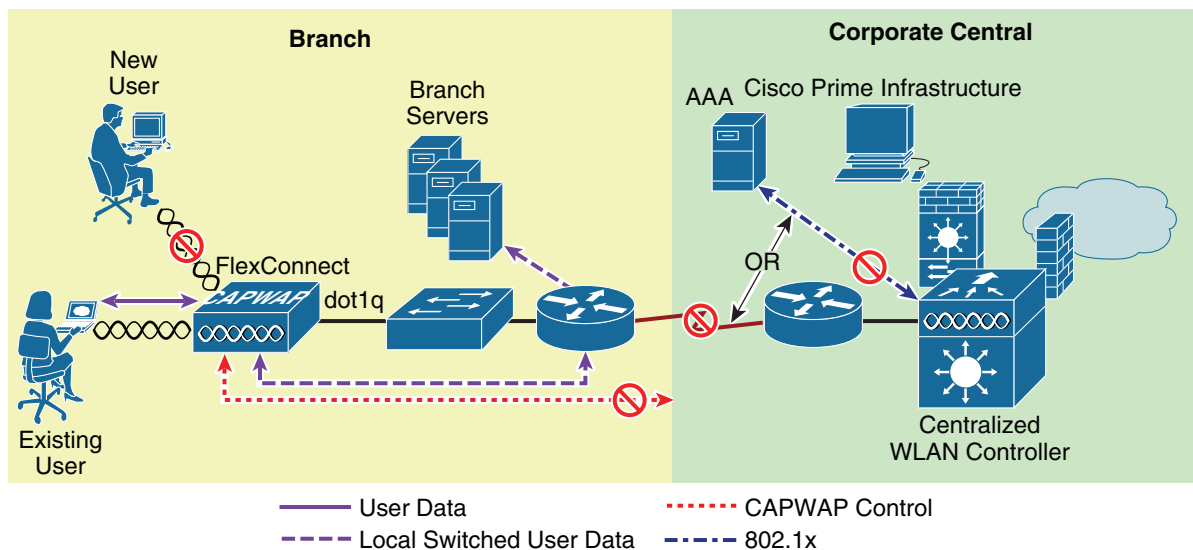
Figure 7-3 Authentication-Central/Switch-Local WLAN



Authentication-Down/Switch-Local

A WLAN that requires central authentication (as explained above) rejects new users. Existing authenticated users continue to be switched locally until session time-out (if configured). The WLAN continues to beacon and respond to probes until there are no more (existing) users associated to the WLAN. This state occurs as a result of the AP going into standalone mode (Figure 7-4).

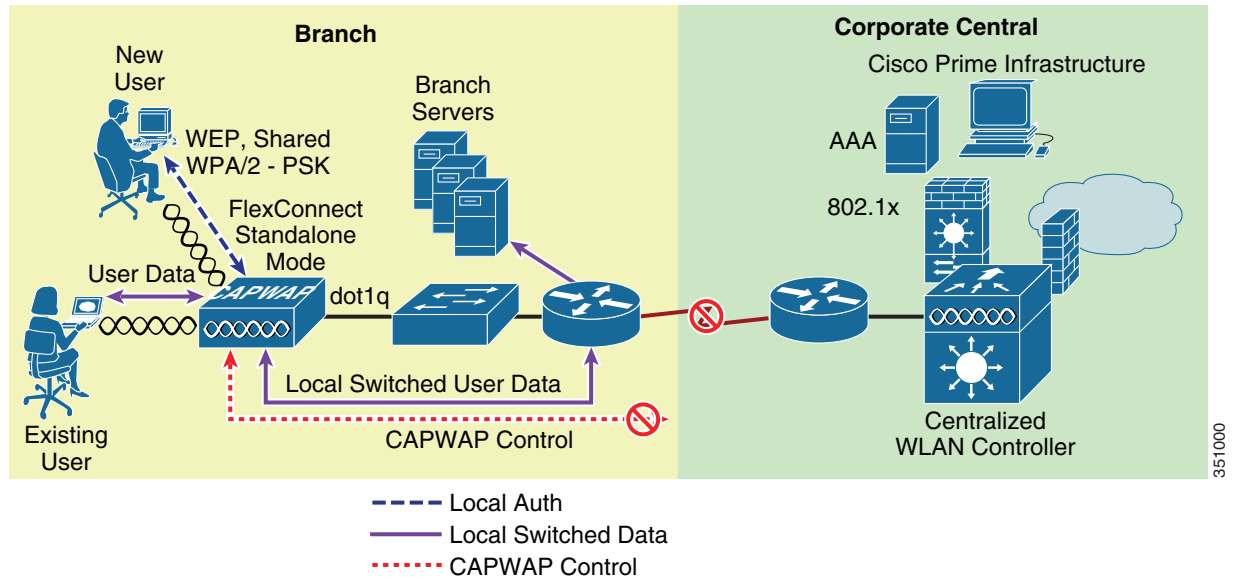
Figure 7-4 Authentication-Down/Local Switch



Authentication-local/switch-local

This state represents a WLAN that uses open, static WEP, shared, or WPA2 PSK security methods. User traffic is switched locally. These are the only security methods supported locally if a FlexConnect goes into standalone mode. The WLAN continues to beacon and respond to probes (Figure 7-5). Existing users remain connected and new user associations are accepted. If the AP is in connected mode, authentication information for these security types is forwarded to the WLC.

Figure 7-5 Authentication-Local/Switch-Local WLAN



Note

All 802.11 authentication and association processing occurs regardless of which operational mode the AP is in. When in connected mode, the FlexConnect AP forwards all association/authentication information to the WLC. When in standalone mode, the AP cannot notify the WLC of such events, which is why WLANs that make use of central authentication/switching methods are unavailable.

Applications

The FlexConnect AP offers greater flexibility in how it can be deployed, such as:

- Branch wireless connectivity
- Branch guest access
- Public WLAN hotspot

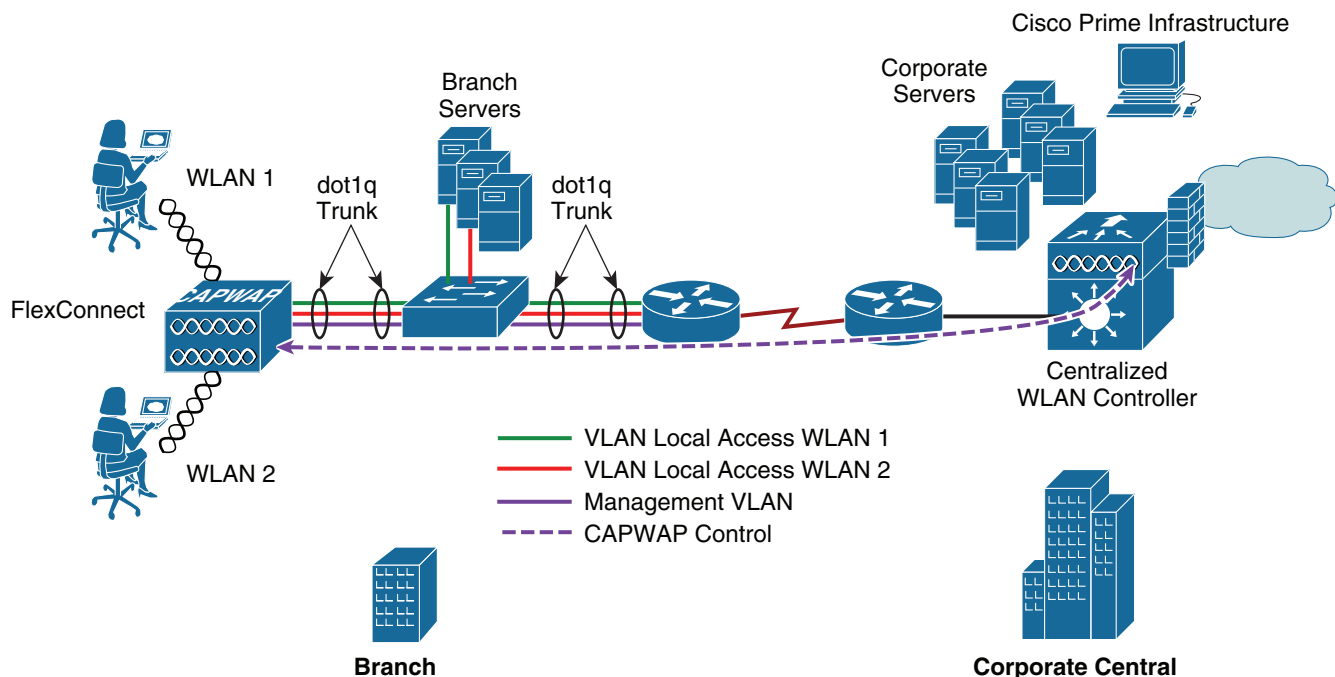
Branch Wireless Connectivity

FlexConnect addresses the wireless connectivity needs in branch locations by permitting wireless user traffic to terminate locally rather than tunneled across the WAN to a central WLC. With FlexConnect, branch locations can more effectively implement segmentation, access control, and QoS policies on a per-WLAN basis, as shown in Figure 7-6.

Branch Guest Access

The centralized WLC itself, as shown in [Figure 7-6](#), can perform web authentication for guest access WLANs. The guest user's traffic is segmented (isolated) from other branch office traffic. For more detailed information on guest access, refer to [Chapter 10, “Cisco Unified Wireless Network Guest Access Services.”](#)

Figure 7-6 FlexConnect Topology

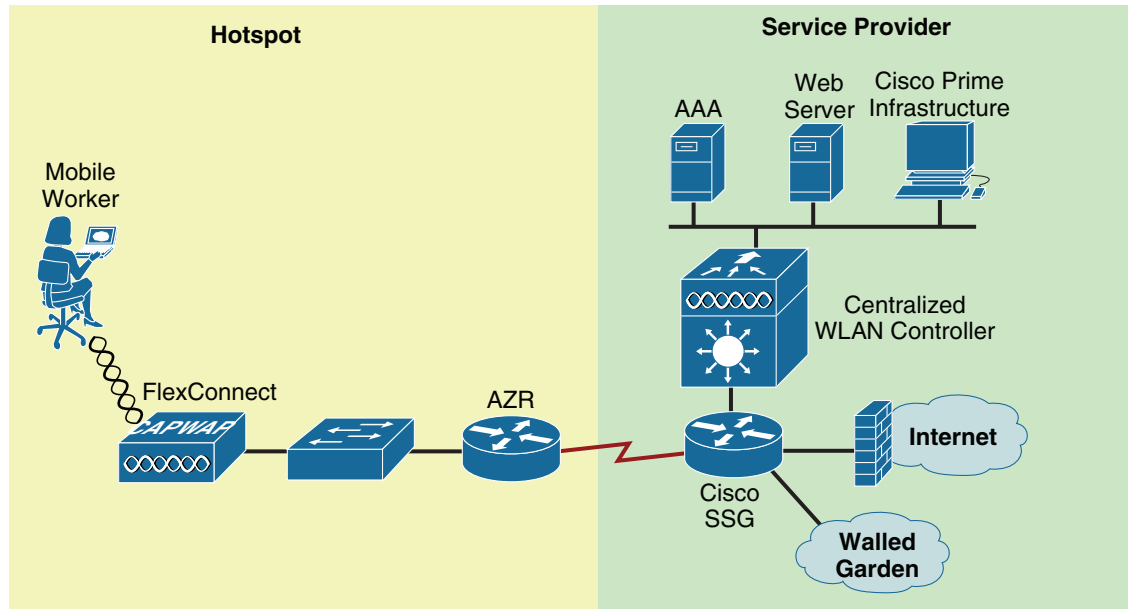


351021

Public WLAN Hotspot

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

The FlexConnect AP, with its ability to map WLANs to separate VLANs, is an alternative to a standalone AP for small venue hotspot deployments where only one, or possibly two, APs are needed. [Figure 7-7](#) provides an example of hotspot topology using a FlexConnect AP.

Figure 7-7 Hotspot Access using FlexConnect Local Switching

351002

Deployment Considerations

The following section covers the various implementation and operational caveats associated with deploying FlexConnect APs.

WAN Link

For the FlexConnect AP to function predictably, keep in mind the following with respect to WAN link characteristics:

- **Latency**—A given WAN link should not impose latencies greater than 100 ms. The AP sends heartbeat messages to the WLC once every thirty seconds. If a heartbeat response is missed, the AP sends five successive heartbeats (one per second) to determine whether connectivity still exists. If connectivity is lost, the FlexConnect AP switches to standalone mode (see [Operation Modes, page 7-3](#) for operation mode definitions). The AP itself is relatively delay tolerant. However, at the client, timers associated with authentication are sensitive to link delay, and thus a constraint of ≤ 100 ms is required. Otherwise, the client can time-out waiting to authenticate, which can cause other unpredictable behaviors, such as looping.
- **Bandwidth**—WAN links should be at least 128 kbps for deployments when up to eight APs are being deployed at a given location. If more than eight APs are deployed, proportionally more bandwidth should be provisioned for the WAN link.
- **Path MTU**—An MTU no smaller than 500 bytes is required.

Roaming

When a FlexConnect AP is in connected mode, all client probes, association requests, 802.1x authentication requests, and corresponding response messages are exchanged between the AP and the WLC via the CAPWAP control plane. This is true for open, static WEP, and WPA PSK-based WLANs even though CAPWAP connectivity is not required to use these authentication methods when the AP is in standalone mode.

- **Dynamic WEP/WPA**—A client that roams between FlexConnect APs using one of these key management methods performs full authentication each time it roams. After successful authentication, new keys are passed back to the AP and client. This behavior is no different than a standard centralized WLAN deployment, except that in an FlexConnect topology, there can be link delay variations across the WAN, which can in turn impact total roam time. Depending on the WAN characteristics, RF design, back end authentication network, and authentication protocols being used, roam times may vary.
- **WPA2**—To improve client roam times, WPA2 introduced key caching capabilities, based on the IEEE 802.11i specification. Cisco created an extension to this specification called Proactive Key Caching (PKC). PKC today is supported only by the Microsoft Zero Config Wireless supplicant and the Funk (Juniper) Odyssey client. Cisco CCKM is also compatible with WPA2.

FlexConnect does not support PKC, regardless of whether a WLAN is centrally or locally-switched. As such, PKC-capable clients that roam between FlexConnect APs undergo full 802.1x authentication. Remote branch locations requiring predictable, fast roaming behavior in support of applications such as wireless IP telephony should consider deploying a local WLC (Virtual Controller on UCS blade or 2500 WLC).

- **Cisco Centralized Key Management (CCKM)**—CCKM is a Cisco-developed protocol in which the WLC caches the security credentials of CCKM-capable clients and forwards those credentials to other APs within a mobility group. When a client roams and associates with another AP, their credentials are forwarded to that AP, which allows the client to re-associate and authenticate in a two-step process. This eliminates the need for full authentication back to the AAA server. CCKM-capable clients undergo full 802.1x authentication each time they roam from one FlexConnect to another.
- **Layer 2 switch CAM table updates**—When a client roams from one AP to another on a locally-switched WLAN, FlexConnect does not announce to a Layer 2 switch that the client has changed ports. The switch will not discover that the client has roamed until the client performs an ARP request for its default router. This behavior, while subtle, can have an impact on roaming performance.



Note

A client that roams (for a given local switched WLAN) between FlexConnect APs that map the WLAN to a different VLAN/subnet will renew their IP addresses to ensure that they have an appropriate address for the network to which they have roamed.

Radio Resource Management

While in connected mode, all radio resource management (RRM) functionality is fundamentally available. However, because typical FlexConnect deployments comprise a smaller number of APs, RRM functionality might not be operational at a branch location. For example, in order for transmit power control (TPC) to work, there must be a minimum of four FlexConnect APs in proximity to each other. Without TPC, other features such as coverage hole protection will be unavailable.

Location Services

FlexConnect deployments typically consist of only a handful of APs at a given location. Cisco maintains strict guidelines regarding the number and placement of APs to achieve the highest level of location accuracy. As such, although it is possible to obtain location information from FlexConnect deployments, the level of accuracy may vary greatly across remote location deployments.

QoS Considerations

For WLANs that are centrally-switched, the FlexConnect AP handles QoS in the same way as standard APs. Locally-switched WLANs implement QoS differently.

For locally-switched WLANs with Wi-Fi MultiMedia (WMM) traffic, the AP marks the dot1p value within the dot1q VLAN tag for upstream traffic. This happens only for tagged VLANs, not the native VLAN.

For downstream traffic, FlexConnect uses the incoming dot1p tag from the locally-switched Ethernet and uses this to queue and mark the WMM values associated with frames destined to a given user across the RF link.

The WLAN QoS profile is applied both for upstream and downstream packets. For downstream, if an 802.1p value that is higher than the default WLAN value is received, the default WLAN value is used. For upstream, if the client sends a WMM value that is higher than the default WLAN value, the default WLAN value is used. For non-WMM traffic, there is no CoS marking on the client frames from the AP.

For more information see [Chapter 5, “Cisco Unified Wireless QoS.”](#)

**Note**

Cisco strongly recommends that appropriate queuing/policing mechanisms be implemented across the WAN to ensure proper handling of traffic based on its DSCP setting. An appropriate priority queue should be reserved for CAPWAP control traffic to ensure that a FlexConnect AP does not inadvertently cycle between connected and standalone modes because of congestion.

General Deployment Considerations

Although it is possible for any WLC to support FlexConnect APs, depending on the number of branch locations and subsequently the total number of APs being deployed, it makes sense (from an administrative standpoint) to consider using a dedicated WLC(s) to support a FlexConnect deployment.

FlexConnect APs typically do not share the same policies as APs within a main campus; each branch location is essentially an RF and mobility domain unto itself. Even though a single WLC cannot be partitioned into multiple logical RF and mobility domains, a dedicated WLC allows branch-specific configuration and policies to be logically separate from the campus.

If deployed, a dedicated FlexConnect WLC should be configured with a different mobility and RF network name than that of the main campus. All FlexConnect APs joined to the *dedicated* WLC become members of that RF and mobility domain.

From an auto-RF standpoint, assuming there are enough FlexConnect APs deployed within a given branch (see [Radio Resource Management, page 7-8](#)), the WLC attempts to auto manage the RF coverage associated with each branch.

There is no advantage (or disadvantage) in having the FlexConnect APs consolidated into their own mobility domain. This is because client traffic is switched locally. EoIP mobility tunnels are not invoked between WLCs (of the same mobility domain) where client roaming with FlexConnect APs is involved.

If a dedicated WLC is going to be used for a FlexConnect deployment, a backup WLC should also be deployed to ensure network availability. As with standard AP deployments, the WLC priority should be set on the FlexConnect APs to force association with the designated WLCs.

FlexConnect Solution

The FlexConnect solution enables you to:

- Centralize control and management traffic.
- Distribute the client data traffic at each Branch Office.
- Ensure traffic flow is going to its final destination in the most efficient manner.

Advantages of Centralizing Access Point Control Traffic

The advantages of centralizing AP control traffic are:

- Single pane of monitoring and troubleshooting
- Ease of management
- Secured and seamless mobile access to Data Center resources
- Reduction in branch footprint
- Increase in operational savings

Advantages of Distributing Client Data Traffic

The advantages of distributing client data traffic are:

- No operational downtime (survivability) against complete WAN link failures or controller unavailability.
- Mobility resiliency within branch during WAN link failures.
- Increase in branch scalability. Supports branch size that can scale up to 100 APs and 250,000 square feet (5000 square feet per AP).

Central Client Data Traffic

The Cisco FlexConnect solution also supports Central Client Data Traffic, but it should be limited to Guest data traffic only. [Table 7-1](#) and [Table 7-2](#) outline the restrictions on WLAN security types only for non-guest clients whose data traffic is also switched centrally at the Data Center.

Table 7-1 *Layer 2 Security Support for Centrally-Switched Non-Guest Users*

WLAN Layer 2 Security	Type	Results
None	N/A	Allowed
WPA + WPA2	802.1x	Allowed
	CCKM	Allowed
	802.1x + CCKM	Allowed
	PSK	Allowed
802.1x	WEP	Allowed
Static WEP	WEP	Allowed
WEP + 802.1x	WEP	Allowed
CKIP		Allowed

**Note**

These authentication restrictions do not apply to clients whose data traffic is distributed at the branch.

Table 7-2 *Layer 3 Security Support for Centrally and Locally Switched Users*

WLAN Layer 3 Security	Type	Results
Web Authentication	Internal	Allowed
	External	Allowed
	Customized	Allowed
Web Pass-Through	Internal	Allowed
	External	Allowed
	Customized	Allowed
Conditional Web Redirect	External	Allowed
Splash Page Redirect	External	Allowed

Cisco Flex 7500 Series Cloud Controller

The Cisco Flex 7500 Series Cloud Controller can manage wireless APs in up to 500 branch locations allowing IT managers to configure, manage, and troubleshoot up to 3000 APs and 30,000 clients from the data center. The Cisco Flex 7500 Cloud Series Controller supports secure guest access, rogue detection for Payment Card Industry (PCI) compliance, and in-branch (locally-switched) Wi-Fi voice and video.

The Cisco Flex 7500 Series Cloud Controller supports the following APs: 1040, 1130, 1140, 1240, 1250, 1260, 1550, , 2600, 3500, 3600, OEAP 600, ISR 881, and ISR 891. These APs support multiple SSIDs.

Table 7-3 compares the scalability between the Flex 7500, WiSM-2 and WLC 5500 Series controllers.

Table 7-3 Controller Scalability Comparison

Scalability	Flex 7500	WiSM-2	WLC 5500
Total Access Points	6,000	1000	500
Total Clients	64,000	15,000	7,000
Max FlexConnect Groups	2,000	100	100
Max APs per FlexConnect Group	100	25	25
Max AP Group	600	1000	500

Modes of Operation

There are two modes of operation for FlexConnect. They are:

Connected—FlexConnect is said to be in connected mode when its CAPWAP control plane, back to the controller, is up and operational. That is, the WAN link is up and functioning as expected.

Standalone—FlexConnect enters into standalone mode when it no longer has connectivity back to the controller. FlexConnect APs in standalone mode will continue to function, with the last known configuration, even in the event of power failure or WAN failure.

Primary Design Requirements

FlexConnect APs are deployed at the Branch site and managed from the Data Center over a WAN link. It is highly recommended that the minimum bandwidth restriction remains 12.8 kbps per AP with the round trip latency no greater than 300 ms for data deployments and 100 ms for data + voice deployments (see Table 7-4). The maximum transmission unit (MTU) must be at least 500 bytes.

Table 7-4 Bandwidth Minimums

Deployment Type	WAN Bandwidth (Min)	WAN RTT Latency (Max)	APs per Branch (Max)	Clients per Branch (Max)
Data	64 kbps	300 ms	5	25
Data + Voice	128 kbps	100 ms	5	25
Monitor	64 kbps	2 sec	5	N/A
Data	640 kbps	300 ms	50	1000
Data + Voice	1.44 Mbps	100 ms	50	1000
Monitor	640 kbps	2 sec	50	N/A

The primary design requirements are:


- Branch size that can scale up to 100 APs and 250,000 square feet (5000 square feet per AP)
- Central management and troubleshooting
- No operational downtime

- Client-based traffic segmentation
- Seamless and secured wireless connectivity to corporate resources
- PCI compliant
- Support for guests

Branch Networking Features and Best Practices

The FlexConnect solution virtualizes the complex security, management, configurations and troubleshooting operations within the data center and then transparently extends those services to each branch. Deployments using FlexConnect controllers are easier for IT to set up, manage and, most importantly, scale (see [Figure 7-8](#)).

Figure 7-8 *Wireless Branch Network Design*



The features and best practices include:

- FlexConnect Groups—Provides the functionality of Local Backup Radius, Cisco's Centralized Key Management (CCKM) and Opportunistic Key Caching (OKC) fast roaming and Local Authentication.
- Fault Tolerance—Improves the wireless branch resiliency without operational downtime.
- ELM (Enhanced Local Mode for Adaptive wIPS)—Provides adaptive wIPS functionality when serving clients without any impact to client performance.
- Client Limit per WLAN—Limiting total guest clients on branch network.
- Auto-convert APs in FlexConnect—Functionality to automatically convert APs in FlexConnect for the branch.
- Guest Access—Continue existing Cisco's Guest Access Architecture with FlexConnect.

FlexConnect Groups

Because all of the FlexConnect APs at each branch site are part of a single FlexConnect Group, FlexConnect Groups ease the organization of each branch site.



Note

FlexConnect Groups are not analogous to AP Groups.

The FlexConnect Group is primarily designed to solve the following challenges:

- How can wireless clients perform 802.1X authentication and access Data Center services if the controller fails?
- How can wireless clients perform 802.1X authentication if WAN link between Branch and Data Center fails?
- Is there any impact on branch mobility during WAN failures?
- Does the FlexConnect Solution provide no operational branch downtime?

You can configure the controller to allow a FlexConnect AP, in standalone mode, to perform full 802.1X authentication to a backup RADIUS server.



Note

Backup RADIUS accounting is not supported.

In order to increase the resiliency of the branch, administrators can configure a primary backup RADIUS server or both a primary and secondary backup RADIUS server. These servers are used only when the FlexConnect AP is not connected to the controller.

Configuring FlexConnect Groups

Complete the following procedure to configure FlexConnect groups to support Local Authentication using Local Extensible Authentication Protocol (LEAP), when FlexConnect is either in connected or standalone mode.

-
- Step 1** Click New under **Wireless > FlexConnect Groups**.
 - Step 2** Assign Group Name Store 1 (similar to the configuration in [Figure 7-8](#)).
 - Step 3** Click Apply when the Group Name is set.
 - Step 4** Click the newly created Group Name Store 1.
 - Step 5** Click Add AP.
 - Step 6** Check the Enable AP Local Authentication box in order to enable Local Authentication when the AP is in standalone mode.
 - Step 7** Check the Select APs from current controller box in order to enable the AP Name drop-down menu.
 - Step 8** Choose the AP from the drop-down that needs to be part of this FlexConnect Group.

Step 9 Click Add after the AP is chosen from the drop-down.

Step 10 Repeat steps 7 and 8 to add all of the APs to this FlexConnect group Store 1.



Note Maintaining 1:1 ratio between the AP-Group and FlexConnect group simplifies network management.

Step 11 Click the **Local Authentication** tab then the **Protocols** tab and check the Enable LEAP Authentication box.

Step 12 Click Apply after the check box is set.



Note If you have a backup controller, make sure the FlexConnect groups are identical and AP MAC address entries are included per FlexConnect group.

Step 13 Under Local Authentication, click Local Users.

Step 14 Set the Username, Password and Confirm Password fields, then click Add to create user entry in the LEAP server residing on the AP.

Step 15 Repeat step 13 until your local username list is exhausted. You cannot configure or add more than 100 users.

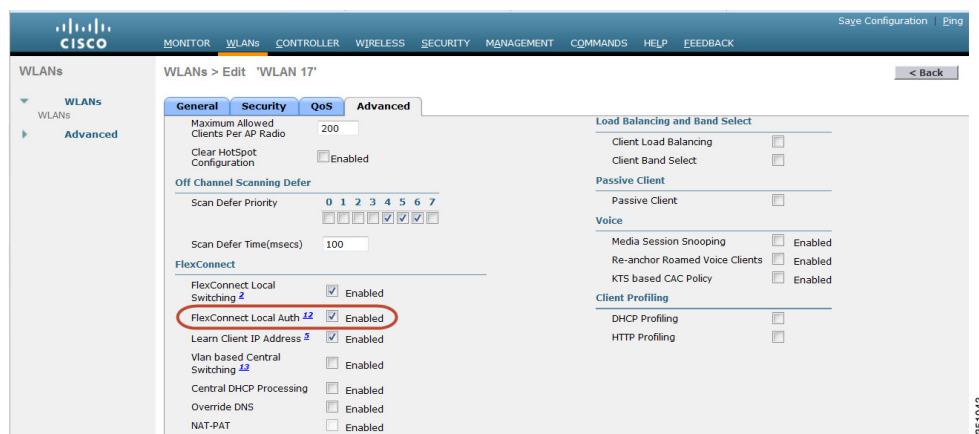
Step 16 Click Apply after entering all local user information. The user count is verified.

Step 17 From the top pane, click WLANs.

Step 18 Click WLAN ID number that was created during the AP Group creation. In this example, WLAN 17

Step 19 Under WLAN > Edit for WLAN ID 17, click Advanced.

Step 20 Check the FlexConnect Local Auth box in order to enable Local Authentication in connected mode.

**Note**

Local Authentication is supported only for FlexConnect with Local Switching. Always make sure to create the FlexConnect Group before enabling Local Authentication under WLAN

CLI Verification

Client authentication state and switching mode can quickly be verified using this CLI command on the WLC:


```
(Cisco Controller) >show client detail 00:24:d7:2b:7c:0c
Client MAC Address..... 00:24:d7:2b:7c:0c
Client Username..... N/A
AP MAC Address..... d0:57:4c:08:e6:70
Client State..... Associated
FlexConnect Data Switching..... Local
FlexConnect Authentication..... Local
```

Local Authentication

Figure 7-9 illustrates clients continuing to perform 802.1X authentication even after the FlexConnect Branch APs lose connectivity with the controller. As long as the RADIUS/ACS server is reachable from the Branch site, wireless clients will continue to authenticate and access wireless services.

In other words, if the RADIUS/ACS is located inside the Branch, then clients will authenticate and access wireless services even during a WAN outage.

Figure 7-9 Local Authentication—AP Authenticator



Note

This feature can be used in conjunction with the FlexConnect backup RADIUS server feature. If a FlexConnect Group is configured with both backup RADIUS server and local authentication, the FlexConnect AP always attempts to authenticate clients using the primary backup RADIUS server first, followed by the secondary backup RADIUS server (if the primary is not reachable), and finally, the Local EAP Server on FlexConnect AP itself (if the primary and secondary are not reachable).

Local EAP

You can configure the controller to allow a FlexConnect AP in standalone or connected mode to perform LEAP or EAP-FAST authentication for up to 100 statically configured users. The controller sends the static list of user names and passwords to each FlexConnect AP of that particular FlexConnect Group when it joins the controller. Each AP in the group authenticates its own associated clients.

This feature is ideal for customers who are migrating from a standalone AP network to a lightweight FlexConnect AP network and are *not* interested in maintaining a large user database or adding another hardware device to replace the RADIUS server functionality available in the standalone AP.

CCKM/OKC Fast Roaming

FlexConnect Groups are required for Cisco's Centralized Key Management (CCKM) and Opportunistic Key Caching (OKC) fast roaming to work with FlexConnect APs. Fast roaming is achieved by caching a derivative of the master key from a full EAP authentication so that a simple and secure key exchange can occur when a wireless client roams to a different AP.

This feature prevents the need to perform a full RADIUS EAP authentication as the client roams from one AP to another. The FlexConnect APs need to obtain the CCKM/OKC cache information for all the clients that might associate so they can process it quickly instead of sending it back to the controller.

For example, if you have a controller with 300 APs and 100 clients that might associate, sending the CCKM/OKC cache for all 100 clients may not be practical. If you create a FlexConnect Group comprising a limited number of APs (for example, you create a group for four APs in a remote office), the clients will then roam only among those four APs, and the CCKM/OKC cache is distributed among those four APs only when the clients associate to one of them.

This feature along with backup RADIUS and Local Authentication (Local-EAP) ensures no operational downtime for your branch sites.



Note

CCKM/OKC fast roaming is supported on FlexConnect APs only.

FlexConnect VLAN Override

In the current FlexConnect architecture, there is a strict mapping of WLAN to VLAN, and thus the client getting associated on a particular WLAN on a FlexConnect AP has to abide by a VLAN that is mapped to it. This method has limitations because it requires clients to associate with different SSIDs in order to inherit different VLAN-based policies.

From 7.2 release onwards, AAA override of VLAN on individual WLAN configured for local switching is supported. In order to have a dynamic VLAN assignment, APs would have the interfaces for the VLAN pre-created based on a configuration using existing WLAN-VLAN mapping for individual FlexConnect APs or using ACL-VLAN mapping on a FlexConnect group. The WLC is used to pre-create the sub-interfaces at the AP.

FlexConnect VLAN Override Summary

- AAA VLAN override is supported from release 7.2 for WLANs configured for local switching in central and local authentication mode.
- AAA override should be enabled on WLANs configured for local switching.
- The FlexConnect AP should have VLAN pre-created from WLC for dynamic VLAN assignment.
- If VLANs returned by AAA override are not present on AP clients, they will get an IP from the default VLAN interface of the AP.

FlexConnect VLAN Based Central Switching

From release 7.3 onwards, traffic from FlexConnect APs can be switched centrally or locally depending on the presence of a VLAN on a FlexConnect AP.

In controller software release 7.2, AAA override of VLAN (Dynamic VLAN assignment) for locally-switched WLANs puts wireless clients on the VLAN provided by the AAA server. If the VLAN provided by the AAA server is not present at the AP, the client is put on a WLAN mapped VLAN on that AP and traffic switches locally on that VLAN. Further, prior to release 7.3, traffic for a particular WLAN from FlexConnect APs can be switched Centrally or Locally depending on the WLAN configuration.

FlexConnect VLAN Central Switching Summary

Traffic flow on WLANs configured for Local Switching when FlexConnect APs are in connected mode are as follows:

- If the VLAN is returned as one of the AAA attributes and that VLAN is not present in the FlexConnect AP database, traffic will switch centrally and the client is assigned this VLAN/Interface returned from the AAA server provided that the VLAN exists on the WLC.
- If the VLAN is returned as one of the AAA attributes and that VLAN is not present in the FlexConnect AP database, traffic will switch centrally. If that VLAN is also not present on the WLC, the client will be assigned a VLAN/Interface mapped to a WLAN on the WLC.
- If the VLAN is returned as one of the AAA attributes and that VLAN is present in the FlexConnect AP database, traffic will switch locally.
- If the VLAN is not returned from the AAA server, the client is assigned a WLAN mapped VLAN on that FlexConnect AP and traffic is switched locally.

Traffic flow on WLANs configured for Local Switching when FlexConnect APs are in standalone mode are as follows:

- If the VLAN returned by the AAA server is not present in the FlexConnect AP database, the client will be put on a default VLAN (that is, a WLAN mapped VLAN on a FlexConnect AP). When the AP connects back, this client is de-authenticated and will switch traffic centrally.
- If the VLAN returned by the AAA server is present in the FlexConnect AP database, the client is placed into a returned VLAN and traffic will switch locally.
- If the VLAN is not returned from the AAA server, the client is assigned a WLAN mapped VLAN on that FlexConnect AP and traffic will switch locally.

FlexConnect ACL

With the introduction of ACLs on FlexConnect, there is a mechanism to cater to the need of access control at the FlexConnect AP for protection and integrity of locally-switched data traffic from the AP. FlexConnect ACLs are created on the WLC and should then be configured with the VLAN present on the FlexConnect AP or FlexConnect group using VLAN-ACL mapping, which will be for AAA override VLANs. These are then pushed to the AP.

FlexConnect ACL Summary

- Create FlexConnect ACL on the controller.
- Apply the same on a VLAN present on FlexConnect AP under AP Level VLAN ACL mapping.
- Can be applied on a VLAN present in FlexConnect Group under VLAN-ACL mapping (generally done for AAA overridden VLANs).
- While applying ACL on VLAN, select the direction to be applied: *ingress*, *egress*, or *ingress and egress*.

FlexConnect ACL Limitations

- A maximum of 512 FlexConnect ACLs can be configured on WLC.

- Each individual ACL can be configured with 64 rules.
- A maximum of 32 ACLs can be mapped per FlexConnect group or per FlexConnect AP.
- At any given point in time, there is a maximum of 16 VLANs and 32 ACLs on the FlexConnect AP.

FlexConnect Split Tunneling

Split Tunneling introduces a mechanism by which the traffic sent by the client will be classified, based on packet content, using FlexConnect ACL. Matching packets are switched locally from FlexConnect AP and the rest of the packets are centrally-switched over CAPWAP.

The Split Tunneling functionality is an added advantage for OEAP setup where clients on a Corporate SSID can talk to devices on a local network (printers, wired machine on a Remote LAN Port, or wireless devices on a Personal SSID) directly without consuming WAN bandwidth by sending packets over CAPWAP.

FlexConnect ACL can be created with rules in order to permit all of the devices present at the local site/network. When packets from a wireless client on the Corporate SSID match the rules in the FlexConnect ACL configured on OEAP, that traffic is switched locally and the rest of the traffic (that is, implicit deny traffic) will switch centrally over CAPWAP.

The Split Tunneling solution assumes that the subnet/VLAN associated with a client in the central site is not present in the local site (that is, traffic for clients that receive an IP address from the subnet present on the central site will not be able to switch locally).

The Split Tunneling functionality is designed to switch traffic locally for subnets that belong to the local site in order to avoid WAN bandwidth consumption. Traffic that matches the FlexConnect ACL rules are switched locally, and NAT operation is performed changing the client's source IP address to the FlexConnect AP's interface IP address that is route-able at the local site/network.

Split Tunnel Summary

- The Split Tunneling functionality is supported on WLANs configured for central switching advertised by FlexConnect APs only.
- The DHCP required should be enabled on WLANs configured for Split Tunneling.
- The Split Tunneling configuration is applied per WLAN configured for central switching on a per FlexConnect AP basis or for all of the FlexConnect APs in a FlexConnect Group.

Split Tunnel Limitations

- FlexConnect ACL rules should not be configured with permit/deny statement with same subnet as source and destination.
- Traffic on a centrally-switched WLAN configured for Split Tunneling can be switched locally only when a wireless client initiates traffic for a host present on the local site. If traffic is initiated by clients/host on a local site for wireless clients on these configured WLANs, the traffic will not be able to reach the destination.
- Split Tunneling is not supported for Multicast/Broadcast traffic. Multicast/Broadcast traffic will switch centrally even if it matches the FlexConnect ACL.

Fault Tolerance

FlexConnect fault tolerance allows wireless access and services to branch clients when the FlexConnect Branch APs:

- Lose connectivity with the primary controller.
- Are switching to the secondary controller.
- Are re-establishing connection to the primary controller.

FlexConnect fault tolerance, along with the local EAP, provides zero branch downtime during a network outage. This feature is enabled by default and cannot be disabled. It requires no configuration on the controller or AP. However, to ensure fault tolerance works smoothly and is applicable, these criteria should be maintained:

- WLAN ordering and configurations have to be identical across the primary and backup controllers.
- VLAN mapping has to be identical across the primary and backup controllers.
- Mobility domain name has to be identical across the primary and backup controllers.
- Use FlexConnect 7500 as both the primary and backup controllers.

Fault Tolerance Summary

- FlexConnect will not disconnect clients when the AP is connecting back to the same controller provided there is no change in configuration on the controller.
- FlexConnect will not disconnect clients when connecting to the backup controller provided there is no change in configuration and the backup controller is identical to the primary controller.
- FlexConnect will not reset its radios on connecting back to the primary controller provided there is no change in configuration on the controller.

Fault Tolerance Limitations

- Supported only for FlexConnect with Central/Local Authentication with Local Switching.
- Centrally-authenticated clients require full re-authentication if the client session timer expires before the FlexConnect AP switches from standalone to connected mode.
- The primary and backup controllers must be in the same mobility domain.

Peer-to-Peer Blocking

Peer-to-peer (P2P) blocking is supported for clients associated on local switching WLAN. Per WLAN, peer-to-peer configuration is pushed by the controller to the FlexConnect AP. P2P blocking can be configured on a WLAN with any of these three actions:

- Disabled—Disables P2P blocking and bridged traffic locally, within the controller, for clients in the same subnet. This is the default value.
- Drop—This causes the controller to discard packets for clients in the same subnet.
- Forward Up-Stream—This forwards a packet on the upstream VLAN. The devices above the controller decide what action to take regarding the packet.

P2P Summary

- P2P Blocking is configured per WLAN
- Per WLAN, P2P blocking configuration is pushed by the WLC to FlexConnect APs.
- P2P blocking action configured as drop or upstream-forward on a WLAN is treated as P2P blocking enabled on the FlexConnect AP.

P2P Limitations

- In FlexConnect solution P2P blocking configuration cannot be applied only to a particular FlexConnect
- AP or subset of APs. It is applied to all FlexConnect APs that broadcast the SSID.
- Unified solution for central switching clients supports P2P upstream-forward. However, this is not supported in the FlexConnect solution. This is treated as P2P drop, and client packets are dropped instead of forwarded to the next network node.
- Unified solution for central switching clients supports P2P blocking for clients associated to different APs. However, this solution targets only clients connected to the same AP. FlexConnect ACLs can be used as a work around for this limitation.

FlexConnect WGB/uWGB Support for Local Switching WLANs

From release 7.3 onward, Cisco's Work Group Bridge/Universal Work Group Bridge (WGB/uWGB) and wired/wireless clients behind WGBs are supported and will work as normal clients on WLANs configured for local switching.

After association, WGB sends the IAPP messages for each of its wired/wireless clients, and FlexConnect APs behave as follows:

- When a FlexConnect AP is in connected mode, it forwards all the IAPP messages to the controller and the controller will process the IAPP messages the same as of local mode AP. Traffic for wired/wireless clients is switched locally from FlexConnect APs.
- An AP in standalone mode processes the IAPP messages; wired/wireless clients on the WGB must be able to register and de-register. Upon transition to connected mode, FlexConnect AP sends the information of wired clients back to the controller. WGB will send registration messages three times when FlexConnect AP transitions from standalone to connected mode.

Wired/Wireless clients will inherit the WGB's configuration, which means no separate configuration like AAA authentication, AAA override, and FlexConnect ACL is required for clients behind WGB.

FlexConnect WGB/uWGB Summary

- No special configuration is required on WLC in order to support WGB on FlexConnect AP.
- Fault Tolerance is supported for WGB and the clients behind WGB.
- WGB is supported on an IOS AP: 1240, 1130, 1140, 1260, and 1250.

FlexConnect WGB/uWGB Limitations

- Wired clients behind WGB will always be on the same VLAN as WGN itself. Multiple VLAN support for clients behind WGB is not supported on the FlexConnect AP for WLANs configured for Local Switching.
- A maximum of 20 clients (wired/wireless) is supported behind WGB when associated to FlexConnect AP on WLAN configured for local switching.
- WebAuth is not supported for clients behind WGB associated on WLANs configured for local switching.

Guidelines and Limitations

- You can deploy a FlexConnect AP with either a static IP address or a DHCP address. A DHCP server must be available locally and must be able to provide the IP address for the AP during boot-up.
- FlexConnect supports up to four fragmented packets or a minimum 500-byte maximum transmission unit (MTU) WAN link.
- CAPWAP control packets must be prioritized over all other traffic.
- Round-trip latency must not exceed 300 milliseconds (ms) between the AP and the controller. If the 300 milliseconds round-trip latency cannot be achieved, configure the AP to perform local authentication.
- FlexConnect includes robust fault tolerance methodology. When the AP and the controller have the same configuration, the connections (rejoin or standby) between the clients and the FlexConnect APs are maintained intact and the clients experience seamless connectivity.
- Client connections are restored only for locally-switched clients that are in the RUN state when the AP moves from standalone mode to connected mode. After the AP moves from the standalone mode to the connected mode, the AP's radio is also reset.
- The primary and secondary controllers for a FlexConnect AP must have the same configuration. Otherwise, the AP might lose its configuration, and certain features (such as WLAN overrides, VLANs, static channel number, and so on) may not operate as expected. In addition, make sure to duplicate the SSID of the FlexConnect AP and its index number on both controllers.
- The controller configuration must *not* change between the time the AP transitions to standalone mode and back to connect mode. Similarly, if the AP is falling back to a secondary or backup controller, the configuration between the primary and secondary or backup controllers must remain the same.
- Session time-out and re-authentication are performed when the AP establishes a connection to the controller.
- If a session timer expires, the client user name, current/support rate, and listen interval values are *reset* to the default values. When the client connection is re-established, the controller does not restore the client's original attributes.
- Multiple FlexConnect groups can be defined in a single location. There is no deployment restriction on the number of FlexConnect APs per location.
- The controller can send multicast packets in the form of unicast or multicast packets to the AP. In FlexConnect mode, the AP can receive multicast packets only in unicast form.
- To use CCKM fast roaming with FlexConnect APs, you must configure FlexConnect Groups.

- FlexConnect APs support a 1-1 network address translation (NAT) configuration and a port address translation (PAT) for all features except true multicast. Multicast is supported across NAT boundaries when configured using the unicast option. FlexConnect APs also support a many-to-one NAT/PAT boundary, except when you want true multicast to operate for all centrally-switched WLANs.

**Note**

Although NAT and PAT are supported for FlexConnect APs, they are not supported on the corresponding controller. Cisco does not support configurations in which the controller is behind a NAT/PAT boundary.

- VPN and PPTP are supported for locally-switched traffic if these security types are accessible locally at the AP.
- NAC out-of-band integration is supported only on WLANs configured for FlexConnect central switching. It is not supported on WLANs configured for FlexConnect *local* switching.
- The QoS per-user bandwidth contracts are supported only on centrally-switched WLANs and APs in the local mode. The QoS profile per-user bandwidth contracts are *not* supported for FlexConnect locally-switched WLANs.
- Guest user configuration is not supported with FlexConnect local switching.
- Workgroup bridges and universal workgroup bridges are supported on FlexConnect APs for locally-switched clients.
- FlexConnect APs do not support client load balancing.
- FlexConnect supports IPv6 clients by bridging the traffic to a local VLAN, similar to IPv4 operation.
- FlexConnect does not support IPv6 ACLs, neighbor discovery caching, or DHCPv6 snooping of IPv6 NDP packets.
- FlexConnect APs with locally-switched WLANs cannot perform IP Source Guard and prevent ARP spoofing. For centrally-switched WLANs, the wireless controller performs the IP Source Guard and ARP Spoofing. To prevent ARP spoofing attacks in FlexConnect APs with local switching, Cisco recommends you use ARP inspection.



Cisco Wireless Mesh Networking

This document provides design and deployment guidelines for the deployment of secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Wireless Mesh Networking solution, a component of the Cisco Unified Wireless Network solution.



Note

For more detailed information about Cisco Wireless Mesh Networking, including configuration and deployment, refer to the *Cisco Mesh Access Points, Design and Deployment Guide, Release 7.3* at <http://www.cisco.com/en/US/docs/wireless/technology/mesh/7.3/design/guide/Mesh.html>.

Mesh networking employs Cisco Aironet 1500 Series outdoor mesh access points (APs) and indoor mesh APs (Cisco Aironet 1040, 1130, 1140, 1240, 1250, 1260, 2600, 3500e, 3500i, 3600e, and 3600i series APs) along with the Cisco Wireless LAN Controller (WLC), and Cisco Prime Infrastructure to provide scalable, central management and mobility between indoor and outdoor deployments. Control and Provisioning of Wireless Access Points (CAPWAP) protocol manages the connection of the mesh APs to the network.

End-to-end security within the mesh network is supported by employing Advanced Encryption Standard (AES) encryption between the wireless mesh APs and Wi-Fi Protected Access 2 (WPA2) clients. This document also outlines radio frequency (RF) components to consider when designing an outdoor network.

The features described in this document are for the following products:

- Cisco Aironet 1550 (1552) Series outdoor 802.11n mesh APs
- Cisco Aironet 1520 (1522, 1524) Series outdoor mesh APs
- Cisco Aironet 1040, 1130, 1140, 1240, 1250, 1260, 2600, 3500e, 3500i, 3600e, and 3600i series indoor mesh APs
- Mesh features in Cisco Wireless LAN Controller
- Mesh features in Cisco Prime Infrastructure



Note

Cisco Aironet 1505 and 1510 mesh APs are not supported because of their End-of-Life status.

Access Point Roles

Access points (APs) within a mesh network operate in one of the following two ways:

- Root AP (RAP)
- Mesh AP (MAP)

MAPs have wireless connections to their controller, while RAPs have wired connections to their controller. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh APs to the controller.

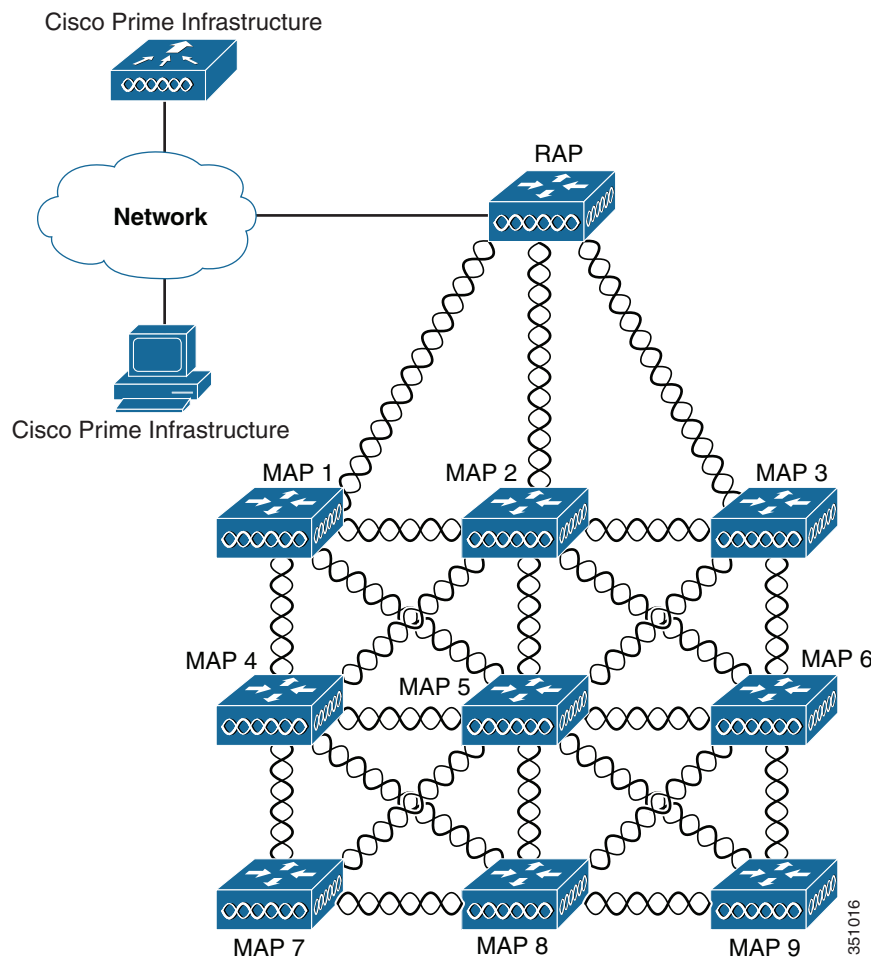


Note

All APs are configured and shipped as mesh APs. To use an AP as a root AP, you must reconfigure the mesh AP to a root AP. In all mesh networks, ensure that there is at least one root AP.

This [Figure 8-1](#) illustrates the relationship between RAPs and MAPs in a mesh network.

Figure 8-1 Simple Mesh Network Hierarchy



Network Access

Wireless mesh networks can simultaneously carry two different traffic types. They are as follows:

- Wireless LAN client traffic
- MAP Ethernet port traffic

Wireless LAN client traffic terminates on the controller, and the Ethernet traffic terminates on the Ethernet ports of the mesh APs.

Access to the wireless LAN mesh for mesh APs is managed by the following authentication methods:

- MAC authentication—Mesh APs are added to a database that can be referenced to ensure they are provided access to a given controller and mesh network.
- External RADIUS Authentication—Mesh APs can be externally authorized using a RADIUS server such as Cisco ACS (4.1 and later) that supports the client authentication type of Identity Services Engine (ISE) or Extensible Authentication Protocol-FAST (EAP-FAST) with certificates.

Network Segmentation

Membership to the wireless LAN mesh network for mesh APs is controlled by the bridge group names (BGNs). Mesh APs can be placed in similar bridge groups to manage membership or provide network segmentation.

Cisco Indoor Mesh Access Points

Indoor mesh is available on the following APs. The list shows the 802.11 protocol supported by each group of APs:

- 802.11a/b/g
 - 1130
 - 1240
- 802.11n
 - 1040
 - 1140
 - 1250
 - 1260
- 802.11n+CleanAir
 - 2600
 - 3500e
 - 3500i
 - 3600

**Note**

For more information about controller software support for APs, see the Cisco Wireless Solutions Software Compatibility Matrix at:

http://www.cisco.com/en/US/docs/wireless/controller/5500/tech_notes/Wireless_Software_Compatibility_Matrix.html.

Enterprise 11n mesh is an enhancement added to the Cisco Unified Wireless Network feature to work with the 802.11n APs. Enterprise 11n mesh features are compatible with non-802.11n mesh but add higher backhaul and client access speeds. The 802.11n indoor APs are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the AP and the other radio can be configured for wireless backhaul. The backhaul is supported only on the 5-GHz radio. Enterprise 11n mesh supports P2P, P2MP, and mesh types of architectures.

You have a choice of ordering indoor APs directly in the bridge mode so that these APs can be used directly as mesh APs. If you have these APs in a local mode (nonmesh), then you have to connect these APs to the controller and change the AP mode to the bridge mode (mesh). This scenario can become cumbersome particularly if the volume of the APs being deployed is large and if the APs are already deployed in the local mode for a traditional nonmesh wireless coverage.

The Cisco indoor mesh APs are equipped with the following two simultaneously operating ratios:

- 2.4-GHz radio used for client access
- 5-GHz radio used for data backhaul

Cisco Outdoor Mesh Access Points

Cisco outdoor mesh APs comprise of the Cisco Aironet 1500 series APs. The 1500 series includes 1552 11n outdoor mesh APs, 1522 dual radio mesh APs, and 1524 multi-radio mesh APs. There are two 1524 models; public safety and 1524PS.

**Note**

In the 6.0 release, the AP1524SB AP was launched in A, C, and N domains. In the 7.0 release, the AP1524SB AP was launched in the -E, -M, -K, -S, and -T domains.

Cisco 1500 series mesh APs are the core components of the wireless mesh deployment. AP1500s are configured by both the controller (GUI and CLI) and Cisco Prime Infrastructure. Communication between outdoor mesh APs (MAPs and RAPs) is over the 802.11a/n radio backhaul. Client traffic is generally transmitted over the 802.11b/g/n radio (802.11a/n can also be configured to accept client traffic), and public safety traffic (AP1524PS only) is transmitted over the 4.9-GHz radio.

The mesh AP can also operate as a relay node for other APs not directly connected to a wired network. Intelligent wireless routing is provided by the Adaptive Wireless Path Protocol (AWPP). This Cisco protocol enables each mesh AP to identify its neighbors and intelligently choose the optimal path to the wired network by calculating the cost of each path in terms of the signal strength and the number of hops required to get to a controller.

AP1500s are manufactured in two different configurations: cable and non-cable.

- The cable configuration can be mounted to a cable strand and supports power-over-cable (POC).
- The non-cable configuration supports multiple antennas. It can be mounted to a pole or building wall and supports several power options.

Uplinks support includes Gigabit Ethernet (1000BASE-T) and a small form-factor (SFP) slot that can be plugged for a fiber or cable modem interface. Both single mode and multimode SFPs up to 1000BASE-BX are supported. The cable modem can be DOCSIS 2.0 or DOCSIS/EuroDOCSIS 3.0 depending upon the type of mesh AP.

AP1500s are available in a hazardous location hardware enclosure. When configured, the AP1500 complies with safety standards for Class I, Division 2, Zone 2 hazardous locations.

The 1520 and 1550 series APs can operate apart from the mesh mode, in the following modes:

- **Local mode**—In this mode, the AP can handle clients on its assigned channel or while monitoring all channels on the band over a 180-second period. During this time, the AP listens on each channel for 50 milliseconds for rogue client beacons, noise floor measurements, interference, and IDS events. The AP also scans for CleanAir interference on the channel.
- **FlexConnect mode**—FlexConnect is a wireless solution for branch office and remote office deployments. The FlexConnect mode enables you to configure and control APs in a branch or remote office from the corporate office through a WAN link without having to deploy a controller in each office. The FlexConnect mode can switch client data traffic locally and perform client authentication locally when the connection to the controller is lost. When connected to the controller, the FlexConnect mode can also tunnel traffic back to the controller.
- **Monitor mode**—In this mode, the AP radios are in the receive state. The AP scans all of the channels every 12 seconds for rogue client beacons, noise floor measurements, interference, IDS events, and CleanAir intruders.
- **Rogue Detector mode**—In this mode, the AP radio is turned off, and the AP listens only to the wired traffic. The controller passes the APs that are configured as rogue detectors as well as lists of suspected rogue clients and AP MAC addresses. The rogue detector listens for ARP packets and can be connected to all broadcast domains through a trunk link.
- **Sniffer mode**—In this mode, the AP captures and forwards all packets on a channel to a remote device that decodes the packets with packet analyzer software such as Wireshark.

Cisco Aironet 1552 Mesh Access Point

The Cisco Aironet 1550 Series Outdoor Mesh Access Point is a modularized wireless outdoor 802.11n AP designed for use in a mesh network. The AP supports point-to-multipoint mesh wireless connectivity and wireless client access simultaneously. The AP can also operate as a relay node for other APs that are not directly connected to a wired network. Intelligent wireless routing is provided by the Adaptive Wireless Path Protocol (AWPP). This enables the AP to identify its neighbors and intelligently choose the optimal path to the wired network by calculating the cost of each path in terms of signal strength and the number of hops required to get to a controller.

The 1550 series APs leverage 802.11n technology with integrated radio and internal/external antennas. The 1552 outdoor platform consists of Multiple Input Multiple Output (MIMO) WLAN radios. It offers 2x3 MIMO with two spatial streams and beamforming, and comes with integrated spectrum intelligence (CleanAir).

CleanAir provides full 11n data rates while detecting, locating, classifying, and mitigating radio frequency (RF) interference to provide the best client experience possible. CleanAir technology on the outdoor 11n platform mitigates Wi-Fi and non-Wi-Fi interference on 2.4-GHz radios.

The 1550 series APs have two radios: 2.4-GHz and 5-GHz MIMO radios. While the 2.4-GHz radios are used primarily for local access, the 5-GHz radios are used for both local access and wireless backhaul in mesh mode.

**Note**

The 2.4-GHz radios cannot be used for backhaul in 1552 APs.

The 2.4-GHz b/g/n radio has the following features:

- Operates in the 2.4-GHz ISM band.
- Supports channels 1-11 in the United States, 1-13 in Europe, and 1-13 in Japan.
- Has two transmitters for 802.11b/g/n operation.
- You can configure the output power for 5 power levels.
- The radio has three receivers that enable maximum-ratio combining (MRC).

The 5-GHz a/n radio has the following feature:

- Operates in the UNII-2 band (5.25 to 5.35 GHz), UNII-2 Extended/ETSI band (5.47 to 5.725 GHz), and the upper ISM band (5.725 to 5.850 GHz).
- Has two transmitters for 802.11a operation.
- Power settings can change depending on the regulatory domain. You can configure the output power for 5 power levels in 3 dB steps.
- The radio has three receivers that enable maximum-ratio combining (MRC).

The 1550 series APs have the following features:

- Supports modularity of the 1520 series and allows flexibility in radio configuration
- Fully interoperable with the 1520 series APs
- Can also interoperate with legacy clients and offers enhanced backhaul performance
- Multicast VideoStream and HotSpot 2.0 are supported when the AP is configured in local mode.
- AP1552 is QoS capable of supporting quality VoWLAN calls.
- Band select, which notifies a connected client to roam from 2.4 GHz to 5 GHz, is supported.
- DTLS support allows AP1552 to encrypt data in all supported AP modes except bridge mode.
- You can enable CleanAir on the 5-GHz radio by navigating to Wireless > Radios > 802.11a > Configure on the controller GUI.

Cisco 1522 Mesh Access Point

The AP1522 mesh AP (part numbers: AIR-LAP1522AG-X-K9, AIR-LAP1522HZ-X-K9, AIR-LAP1522PC-X-K9) includes two radios: a 2.4-GHz and a 4.9- to 5.8-GHz radio. The 2.4-GHz (802.11b/g) radio is for client access and the 5-GHz (802.11a) radio is used as the backhaul. With the 7.0.116.0 release and later releases, 2.4 GHz is available for backhaul. This feature is applicable only to AP1522.

The 5-GHz radio is a 802.11a radio that covers the 4.9- to 5.8-GHz frequency band and is used as a backhaul. It can also be used for client access if the universal client access feature is enabled.

Cisco 1524SB Mesh Access Point

The AP1524SB mesh AP (part number: AIR-LAP1524SB-X-K9) includes three radios: one 2.4-GHz radio and two 5-GHz radios.

The 2.4-GHz radio is for client access (nonpublic safety traffic). The two 5-GHz radios serve as serial backhauls: one uplink and one downlink. The AP1524SB is suitable for linear deployments.



Note

In the 6.0 release, the 5-GHz radios in the –A domain could be operated only in the 5.8-GHz band with 5 channels. In the 7.0 release, these radios cover the whole 5-GHz band.

Each 5-GHz radio backhaul is configured with a different backhaul channel. There is no need to use the same shared wireless medium between the north-bound and south-bound traffic in a mesh tree-based network.

On the RAP, the radio in slot 2 is used to extend the backhaul in the downlink direction; the radio in slot 1 is used only for client access and not mesh.

On the MAP, the radio in slot 2 is used for the backhaul in the uplink direction; the radio in slot 1 is used for the backhaul in the downlink direction.

You only need to configure the RAP downlink (slot 2) channel. The MAPs automatically select their channels from the channel subset. The available channels for the 5.8-GHz band are 149, 153, 157, 161, and 165.

Ethernet Ports

AP1500s support four Gigabit Ethernet interfaces.

- Port 0 (g0) is a Power over Ethernet (PoE) input port–PoE (in)
- Port 1 (g1) is a PoE output port–PoE (out)
- Port 2 (g2) is a cable connection
- Port 3 (g3) is a fiber connection

You can query the status of these four interfaces in the controller CLI and Cisco Prime Infrastructure.

In the controller CLI, the **show mesh env summary** command is used to display the status of the ports. The Up or Down (Dn) status of the four ports is reported in the following format:

```
port0 (PoE-in) : port1 (PoE-out) : port2 (cable) : port3 (fiber)
```

For example, *rap1522.a380* in the display below shows a port status of *UpDnDnDn*. This indicates the following:

PoE-in port 0 (g0) is Up, PoE-out port 1 (g1) is Down (Dn), Cable port 2 (g2) is Down (Dn), and Fiber port 3 (g3) is Down (Dn)

```
(controller)> show mesh env summary
AP Name           Temperature(C/F)  Heater  Ethernet  Battery
-----
rap1242.c9ef      N/A              N/A     UP         N/A
rap1522.a380      29/84            OFF     UpDnDnDn  N/A
rap1522.4da8      31/87            OFF     UpDnDnDn  N/A
```


1550 Series Multiple Power Options

Power options include the following:

Power over Ethernet (PoE)-In

- 56 VDC using a Power Injector (1552E and 1552H)
- PoE-In is not 802.3af and does not work with PoE 802.3af-capable Ethernet switch

AC Power

- 100 to 480 VAC (47-63 Hz)—Connecting AC or Streetlight Power (1552E)
- 100 to 240 VAC—Connecting AC or Streetlight Power (1552H)

External Supply

- 12 VDC—Connecting DC Power Cable (All Models)

Internal Battery Backup (1552E and 1552H)

Power over Cable (PoC)

- 40 to 90VAC—Connecting Cable PoC (1552C)

PoE-Out 802.3af compliant to connect IP devices such as Video Cameras (1552E and 1552H)

- (PoE-Out) is not available when using Power Injector (PoE-In) as the power source

802.3af compliant PoE-Out to connect IP devices such as video cameras (1552E and 1552H)

- This port also performs Auto-MDIX, which enables to connect crossover or straight through cables.

The 1550 series APs can be connected to more than one power source. The APs detect the available power sources and switch to the preferred power source using the following default prioritization:

- AC power or PoC power
- External 12-VDC power
- Power injector PoE power
- Internal battery power

Cisco Wireless LAN Controllers

The wireless mesh solution is supported on Cisco 2500, 5500, and 8500 Series Wireless LAN Controllers. For more information about these controllers, see:

http://www.cisco.com/en/US/products/ps6302/Products_Sub_Category_Home.html.

Cisco Prime Infrastructure

The Cisco Prime Infrastructure provides a graphical platform for wireless mesh planning, configuration, and management. Network managers can use the Cisco Prime Infrastructure to design, control, and monitor wireless mesh networks from a central location.

With the Cisco Prime Infrastructure, network administrators have a solution for RF prediction, policy provisioning, network optimization, troubleshooting, user tracking, security monitoring, and wireless LAN systems management. Graphical interfaces make wireless LAN deployment and operations simple and cost-effective. Detailed trending and analysis reports make the Cisco Prime Infrastructure vital to ongoing network operations.

The Cisco Prime Infrastructure runs on a server platform with an embedded database, which provides scalability that allows hundreds of controllers and thousands of Cisco mesh APs to be managed. Controllers can be located on the same LAN as the Cisco Prime Infrastructure, on separate routed subnets, or across a wide-area connection.

Architecture

Control and Provisioning of Wireless Access Points

Control and Provisioning of Wireless Access Points (CAPWAP) is the protocol used by the controller to manage APs (mesh and nonmesh) in the network. In release 5.2, CAPWAP replaced the lightweight AP protocol (LWAPP).

**Note**

CAPWAP significantly reduces capital expenditures (CapEx) and operational expenses (OpEx), which enables the Cisco wireless mesh networking solution to be a cost-effective and secure deployment option in enterprise, campus, and metropolitan networks.

CAPWAP Discovery on a Mesh Network

The process for CAPWAP discovery on a mesh network is as follows:

1. A mesh AP establishes a link before starting CAPWAP discovery, whereas a nonmesh AP starts CAPWAP discovery using a static IP for the mesh AP, if any.
2. The mesh AP initiates CAPWAP discovery using a static IP for the mesh AP on the Layer 3 network or searches the network for its assigned primary, secondary, or tertiary controller. A maximum of 10 attempts are made to connect.

**Note**

The mesh AP searches a list of controllers configured on the AP (primed) during setup.

3. If Step 2 fails after 10 attempts, the mesh AP falls back to DHCP and attempts to connect in 10 tries.
4. If both Steps 2 and 3 fail and there is no successful CAPWAP connection to a controller, then the mesh AP falls back to LWAPP.
5. If there is no discovery after attempting Steps 2, 3, and 4, the mesh AP tries the next link.

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Traffic Flow

The traffic flow within the wireless mesh can be divided into three components:

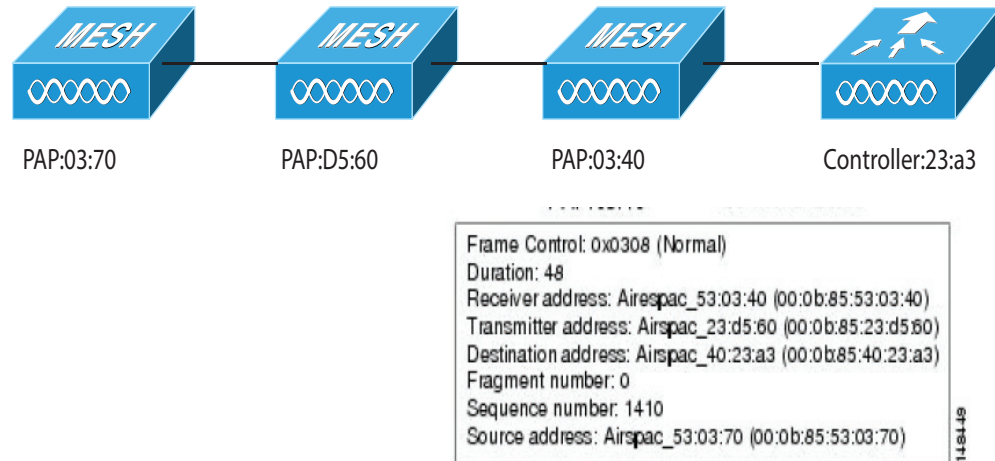
1. Overlay CAPWAP traffic that flows within a standard CAPWAP AP deployment; that is, CAPWAP traffic between the CAPWAP AP and the CAPWAP controller.
2. Wireless mesh data frame flow.
3. AWPP exchanges.

As the CAPWAP model is well known and the AWPP is a proprietary protocol, only the wireless mesh data flow is described. The key to the wireless mesh data flow is the address fields of the 802.11 frames being sent between mesh APs.

An 802.11 data frame can use up to four address fields: receiver, transmitter, destination, and source. The standard frame from a WLAN client to an AP uses only three of these address fields because the transmitter address and the source address are the same. However, in a WLAN bridging network, all four address fields are used because the source of the frame might not be the transmitter of the frame, because the frame might have been generated by a device behind the transmitter.

Figure 8-2 shows an example of this type of framing. The source address of the frame is MAP:03:70, the destination address of this frame is the controller (the mesh network is operating in Layer 2 mode), the transmitter address is MAP:D5:60, and the receiver address is RAP:03:40.

Figure 8-2 Wireless Mesh Frame



As this frame is sent, the transmitter and receiver addresses change on a hop-by-hop basis. AWPP is used to determine the receiver address at each hop. The transmitter address is known because it is the current mesh AP. The source and destination addresses are the same over the entire path.

If the RAP's controller connection is Layer 3, the destination address for the frame is the default gateway MAC address, because the MAP has already encapsulated the CAPWAP in the IP packet to send it to the controller, and is using the standard IP behavior of using ARP to find the MAC address of the default gateway.

Each mesh AP within the mesh forms an CAPWAP session with a controller. WLAN traffic is encapsulated inside CAPWAP and is mapped to a VLAN interface on the controller. Bridged Ethernet traffic can be passed from each Ethernet interface on the mesh network and does not have to be mapped to an interface on the controller.

Mesh Neighbors, Parents, and Children

Relationships among mesh APs are as a parent, child, or neighbor.

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
- Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease values are lower than that of the parent.

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state, and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor, and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.


A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Ease Calculation

Ease is calculated using the SNR and hop value of each neighbor, and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

Figure 8-3 shows the parent path selection where MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.

Figure 8-3 Parent Path Selection



Parent Decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP. That is:

$$\text{adjusted ease} = \min (\text{ease at each hop}) \text{ Hop count}$$

Mesh Deployment Modes

In a Cisco wireless outdoor mesh network, multiple mesh APs comprise a network that provides secure, scalable outdoor wireless LAN.

The three RAPs are connected to the wired network at each location and are located on the building roof. All the downstream APs operate as MAPs and communicate using wireless links (not shown).

Both MAPs and RAPs can provide WLAN client access; however, the location of RAPs are often not suitable for providing client access. All the three APs in are located on the building roofs and are functioning as RAPs. These RAPs are connected to the network at each location.

Some of the buildings have onsite controllers to terminate CAPWAP sessions from the mesh APs but it is not a mandatory requirement because CAPWAP sessions can be back hauled to a controller over a wide-area network (WAN).

Wireless Backhaul

In a Cisco wireless backhaul network, traffic can be bridged between MAPs and RAPs. This traffic can be from wired devices that are being bridged by the wireless mesh or CAPWAP traffic from the mesh APs. This traffic is always AES encrypted when it crosses a wireless mesh link such as a wireless backhaul.

AES encryption is established as part of the mesh AP neighbor relationship with other mesh APs. The encryption keys used between mesh APs are derived during the EAP authentication process.

Only 5 GHz backhaul is possible on all mesh APs except 1522 in which either 2.4 or 5 GHz radio can be configured as a backhaul radio (see Configuring Advanced Features).


Universal Access

You can configure the backhaul on mesh APs to accept client traffic over its 802.11a radio. This feature is identified as Backhaul Client Access in the controller GUI (Monitor > Wireless). When this feature is disabled, backhaul traffic is transmitted only over the 802.11a or 802.11a/n radio and client association is allowed only over the 802.11b/g or 802.11b/g/n radio. For more information about the configuration, see the Configuring Advanced Features.

Point-to-Multipoint Wireless Bridging

In the point-to-multipoint bridging scenario, a RAP acting as a root bridge connects multiple MAPs as nonroot bridges with their associated wired LANs. By default, this feature is disabled for all MAPs. If Ethernet bridging is used, you must enable it on the controller for the respective MAP and for the RAP.

Figure 8-4 shows a simple deployment with one RAP and two MAPs, but this configuration is fundamentally a wireless mesh with no WLAN clients. Client access can still be provided with Ethernet bridging enabled, although if bridging between buildings, MAP coverage from a high rooftop might not be suitable for client access.

Figure 8-4 Point-to-Multipoint Bridging Example

For security reasons the Ethernet port on the MAPs is disabled by default. It can be enabled only by configuring Ethernet bridging on the Root and the respective MAPs. To enable Ethernet bridging using the controller GUI, choose **Wireless > All APs > Details** for the AP page, click the **Mesh** tab, and then select the Ethernet Bridging check box.

Ethernet bridging has to be enabled for the following two scenarios:

- When you want to use the mesh nodes as bridges.
- When you want to connect Ethernet devices such as a video camera on the MAP using its Ethernet port.

Ensure that you enable Ethernet bridging for every parent mesh AP taking the path from the mesh AP in question to the controller. For example, if you enable Ethernet bridging on MAP2 in Hop 2, then you must also enable Ethernet bridging on MAP1 (parent MAP), and on the RAP connecting to the controller.

To configure range parameters for longer links, choose **Wireless > Mesh**. Optimum distance (in feet) should exist between the root AP (RAP) and the farthest mesh AP (MAP). Range from the RAP bridge to the MAP bridge has to be mentioned in feet.

The following global parameter applies to all mesh APs when they join the controller and all existing mesh APs in the network:

- Range: 150 to 132,000 feet
- Default: 12,000 feet

Wireless Backhaul Data Rate

Backhaul is used to create only the wireless connection between the APs. The backhaul interface by default is 802.11a or 802.11a/n depending upon the AP. The rate selection is important for effective use of the available RF spectrum. The rate can also affect the throughput of client devices, and throughput is an important metric used by industry publications to evaluate vendor devices.

Dynamic Rate Adaptation (DRA) introduces a process to estimate optimal transmission rate for packet transmissions. It is important to select rates correctly. If the rate is too high, packet transmissions fail resulting in communication failure. If the rate is too low, the available channel bandwidth is not used, resulting in inferior products, and the potential for catastrophic network congestion and collapse.

Data rates also affect the RF coverage and network performance. Lower data rates, for example 6 Mbps, can extend farther from the AP than can higher data rates, for example 300 Mbps. As a result, the data rate affects cell coverage and consequently the number of APs required. Different data rates are achieved by sending a more redundant signal on the wireless link, allowing data to be easily recovered from noise. The number of symbols sent out for a packet at the 1-Mbps data rate is higher than the number of symbols used for the same packet at 11 Mbps. Therefore, sending data at the lower bit rates takes more time than sending the equivalent data at a higher bit rate, resulting in reduced throughput.

A lower bit rate might allow a greater distance between MAPs, but there are likely to be gaps in the WLAN client coverage, and the capacity of the backhaul network is reduced. An increased bit rate for the backhaul network either requires more MAPs or results in a reduced SNR between MAPs, limiting mesh reliability and interconnection.

ClientLink Technology

Many networks still support a mix of 802.11a/g and 802.11n clients. Because 802.11a/g clients (legacy clients) operate at lower data rates, the older clients can reduce the capacity of the entire network. Cisco's ClientLink technology can help solve problems related to adoption of 802.11n in mixed-client networks by ensuring that 802.11a/g clients operate at the best possible rates, especially when they are near cell boundaries.

Advanced signal processing has been added to the Wi-Fi chipset. Multiple transmit antennas are used to focus transmissions in the direction of the 802.11a/g client, increasing the downlink signal-to-noise ratio and the data rate over range, thereby reducing coverage holes and enhancing the overall system performance. This technology learns the optimum way to combine the signal received from a client and then uses this information to send packets in an optimum way back to the client. This technique is also referred to as Multiple-input multiple-output (MIMO) beamforming, transmit beamforming, or cophasing, and it is the only enterprise-class and service provider-class solution in the market that does not require expensive antenna arrays.

The 802.11n systems take advantage of multipath by sending multiple radio signals simultaneously. Each of these signals, called a spatial stream, is sent from its own antenna using its own transmitter. Because there is some space between these antennas, each signal follows a slightly different path to the receiver, a situation called spatial diversity. The receiver has multiple antennas as well, each with its own radio that independently decodes the arriving signals, and each signal is combined with signals from the other receiver radios. This results in multiple data streams receiving at the same time. This enables a higher throughput than previous 802.11a/g systems, but requires an 802.11n capable client to decipher the signal. Therefore, both AP and client need to support this capability. Due to the complexity of issues, in the first generation of mainstream 802.11n chipsets, neither the AP nor client chipsets implemented 802.11n transmit beamforming. Therefore, the 802.11n standard transmit beamforming will be available eventually, but not until the next generation of chipsets take hold in the market. We intend to lead in this area going forward.

We realized that for the current generation of 802.11n APs, while the second transmit path was being well utilized for 802.11n clients (to implement spatial diversity), it was not being fully used for 802.11a/g clients. In other words, for 802.11a/g clients, some of the capabilities of the extra transmit path was lying idle. In addition, we realized that for many networks, the performance of the installed 802.11a/g client base would be a limiting factor on the network.

To take advantage of this fallow capacity and greatly enhance overall network capacity by bringing 802.11a/g clients up to a higher performance level, we created an innovation in transmit beamforming technology, called ClientLink.

ClientLink uses advanced signal processing techniques and multiple transmit paths to optimize the signal received by 802.11a/g clients in the downlink direction without requiring feedback. Because no special feedback is required, Cisco ClientLink works with all existing 802.11a/g clients.

Cisco ClientLink technology effectively enables the AP to optimize the SNR exactly at the position where the client is placed. ClientLink provides a gain of almost 4 dB in the downlink direction. Improved SNR yields many benefits, such as a reduced number of retries and higher data rates. For example, a client at the edge of the cell that might previously have been capable of receiving packets at 12 Mbps could now receive them at 36 Mbps. Typical measurements of downlink performance with ClientLink show as much as 65 percent greater throughput for 802.11a/g clients. By allowing the Wi-Fi system to operate at higher data rates and with fewer retries, ClientLink increases the overall capacity of the system, which means an efficient use of spectrum resources.

ClientLink in the 1552 APs is based on ClientLink capability available in AP3500s. Therefore, the AP has the ability to beamform well to nearby clients and to update beamforming information on 802.11ACKs. Therefore, even if there is no dedicated uplink traffic, the ClientLink works well, which is beneficial to both TCP and UDP traffic streams. There are no RSSI watermarks, which the client has to cross to take advantage of this beamforming with Cisco 802.11n APs.

ClientLink can beamform to 15 clients at a time. Therefore, the host must select the best 15 if the number of legacy clients exceeds 15 per radio. AP1552 has two radios, which means that up to 30 clients can be beamformed in time domain.

Although ClientLink is applied to legacy OFDM portions of packets, which refers to 11a/g rates (not 11b) for both indoor and outdoor 802.11n APs, there is one difference between ClientLink for indoor 11n and ClientLink for outdoor 11n. For indoor 11n APs, SW limits the affected rates to 24, 36, 48, and 54 Mbps. This is done to avoid clients sticking to a faraway AP in an indoor environment. SW also does not allow ClientLink to work for those rates for 11n clients because the throughput gain is so minimal. However, there is a demonstrable gain for pure legacy clients. For outdoor 11n APs, we do need more coverage. Thus, three more additional legacy data rates lower than 24 Mbps have been added. ClientLink for outdoors is applicable to legacy data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps.

Controller Planning

The following items affect the number of controllers required in a mesh network:

- Mesh APs (RAPs and MAPs) in the network.
- The wired network that connects the RAP and controllers can affect the total number of APs supported in the network. If this network allows the controllers to be equally available to all APs without any impact on WLAN performance, the APs can be evenly distributed across all controllers for maximum efficiency. If this is not the case, and controllers are grouped into various clusters or PoPs, the overall number of APs and coverage are reduced.
- Number of mesh APs (RAPs and MAPs) supported per controller.

For clarity, nonmesh APs are referred to as local APs in this document.

Table 8-1 Mesh AP Support by Controller Model

Controller Model	Local AP Support (nonmesh) ¹	Maximum Possible Mesh AP Support
5508 ²	500	500
2504 ³	50	50
WiSM2	500	500

1. Local AP support is the total number of nonmesh APs supported on the controller model.
2. For 5508, controllers, the number of MAPs is equal to (local AP support - number of RAPs).
3. For 2504, controllers, the number of MAPs is equal to (local AP support - number of RAPs).

**Note**

Mesh is fully supported on Cisco 5508 Controllers. The Base License (LIC-CT508-Base) is sufficient for indoor and outdoor APs (AP152X). The WPlus License (LIC-WPLUS-SW) is merged with the base license. The WPlus License is not required for indoor mesh APs.

Wireless Mesh Network Coverage Considerations

This section provides a summary of items that must be considered for maximum wireless LAN coverage in an urban or suburban area, to adhere to compliance conditions for respective domains.

The following recommendations assume a flat terrain with no obstacles (green field deployment).

We always recommend that you perform a site survey before taking any real estimations for the area and creating a bill of materials.

Cell Planning and Distance

For the Cisco 1520 Series Access Points

The RAP-to-MAP ratio is the starting point. For general planning purposes, the current ratio is 20 MAPs per RAP.

We recommend the following values for cell planning and distance in non-voice networks:

- RAP-to-MAP ratio—Recommended maximum ratio is 20 MAPs per RAP.
- AP-to-AP distance—A spacing of no more than of 2000 feet (609.6 meters) between each mesh AP is recommended. When you extend the mesh network on the backhaul (no client access), use a cell radius of 1000 feet (304.8 meters).
- Hop count—Three to four hops. One square mile in feet (52802), is nine cells and you can cover one square mile with approximately three or four hops.
- For 2.4 GHz, the local access cell size radius is 600 feet (182.88 meters). One cell size is around 1.310×10^6 , so there are 25 cells per square mile.

Collocating Mesh Access Points

The following recommendations provide guidelines to determine the required antenna separation when you collocate AP1500s on the same tower. The recommended minimum separations for antennas, transmit powers, and channel spacing are addressed.

The goal of proper spacing and antenna selection is to provide sufficient isolation by way of antenna radiation pattern, free space path loss, and adjacent or alternate adjacent channel receiver rejection to provide independent operation of the collocated units. The goal is to have negligible throughput degradation due to a CCA hold-off, and negligible receive sensitivity degradation due to a receive noise floor increase.

You must follow antenna proximity requirements, which depend upon the adjacent and alternate adjacent channel usage.

Collocating AP1500s on Adjacent Channels

If two collocated AP1500s operate on adjacent channels such as channel 149 (5745 MHz) and channel 152 (5765 MHz), the minimum vertical separation between the two AP1500s is 40 feet (12.192 meters) (the requirement applies for mesh APs equipped with either 8 dBi omnidirectional or 17 dBi high-gain directional patch antennas).

If two collocated AP1500s operate on channels 1, 6, or 11 (2412 to 2437 MHz) with a 5.5-dBi omnidirectional antenna, then the minimum vertical separation is 8 feet (2.438 meters).

Collocating AP1500s on Alternate Adjacent Channels

If two collocated AP1500s operate on alternate adjacent channels such as channel 149 (5745 MHz) and channel 157 (5785 MHz), the minimum vertical separation between the two AP1500s is 10 feet (3.048 meters) (the requirements applies for mesh APs equipped with either 8-dBi omnidirectional or 17-dBi high-gain directional patch antennas).

If two collocated AP1500s operate on alternate adjacent channels 1 and 11 (2412 MHz and 2462 MHz) with a 5.5-dBi omnidirectional antenna, then the minimum vertical separation is 2 feet (0.609 meters).

In summary, a 5-GHz antenna isolation determines mesh AP spacing requirements and antenna proximity must be followed and is dependent upon the adjacent and alternate adjacent channel usage.

CleanAir

The 1550 series leverages 802.11n technology with integrated radio and internal/external antennas. The 1550 series APs are based on the same chipset as the present CleanAir capable Aironet 3500 APs. In other words, the 1550 series APs are capable of doing CleanAir.

With the 7.3.101.0 release, 2600 series APs can mesh with each other and can also provide CleanAir functionality.

With the 7.2.103.0 release, 3600 series APs can mesh with each other and can also provide CleanAir functionality.

With the 7.0.116.0 release, 3500 series APs can mesh with each other and can also provide CleanAir functionality.

CleanAir in mesh (1552, 2600, 3500 and 3600) can be implemented on the 2.4-GHz radio and provides clients complete 802.11n data rates while detecting, locating, classifying, and mitigating radio frequency (RF) interference. This provides a carrier class management and customer experience and ensures that you have control over the spectrum in the deployed location. CleanAir enabled RRM technology on the outdoor 11n platform detects, quantifies, and mitigates Wi-Fi and non-Wi-Fi interference on 2.4-GHz radios. AP1552 supports CleanAir in 2.4 GHz client access mode.

CleanAir Advisor

If CleanAir is enabled on a backhaul radio, CleanAir Advisor is activated. CleanAir Advisor generates Air Quality Index (AQI) and Interferer Detection Reports (IDR) but the reports are only displayed in the controller. No action is taken through event driven RRM (ED-RRM). CleanAir Advisor is only present on the 5-GHz backhaul radio of APs in bridge mode.

Wireless Mesh Mobility Groups

A mobility group allows controllers to peer with each other to support seamless roaming across controller boundaries. APs learn the IP addresses of the other members of the mobility group after the CAPWAP Join process. A controller can be a member of a single mobility group which can contain up to 24 controllers. Mobility is supported across 72 controllers. There can be up to 72 members (WLCs) in the mobility list with up to 24 members in the same mobility group (or domain) participating in client hand-offs. The IP address of a client does not have to be renewed in the same mobility domain. Renewing the IP address is irrelevant in the controller-based architecture when you use this feature.

Multiple Controllers

The consideration in distance of the CAPWAP controllers from other CAPWAP controllers in the mobility group, and the distance of the CAPWAP controllers from the RAP, is similar to the consideration of an CAPWAP WLAN deployment in an enterprise.

There are operational advantages to centralizing CAPWAP controllers, and these advantages need to be traded off against the speed and capacity of the links to the CAPWAP APs and the traffic profile of the WLAN clients using these mesh APs.

If the WLAN client traffic is expected to be focused on particular sites, such as the Internet or a data center, centralizing the controllers at the same sites as these traffic focal points gives the operational advantages without sacrificing traffic efficiency.

If the WLAN client traffic is more peer-to-peer, a distributed controller model might be a better fit. It is likely that a majority of the WLAN traffic are clients in the area, with a smaller amount of traffic going to other locations. Given that many peer-to-peer applications can be sensitive to delay and packet loss, you should ensure that traffic between peers takes the most efficient path.

Given that most deployments see a mix of client-server traffic and peer-to-peer traffic, it is likely that a hybrid model of CAPWAP controller placement is used, where points of presence (PoPs) are created with clusters of controllers placed in strategic locations in the network.

The CAPWAP model used in the wireless mesh network is designed for campus networks; that is, it expects a high-speed, low-latency network between the CAPWAP mesh APs and the CAPWAP controller.

Increasing Mesh Availability

In the Cell Planning Distance section, a wireless mesh cell of one square mile was created and then built upon. This wireless mesh cell has similar properties to the cells used to create a cellular phone network because the smaller cells (rather than the defined maximum cell size) can be created to cover the same physical area, providing greater availability or capacity. This process is done by adding a RAP to the

cell. Similar to the larger mesh deployment, the decision is whether to use RAP on the same channel, as shown in [Figure 8-5](#) or to use RAPs placed on different channels, as shown in [Figure 8-6](#). The addition of RAPs into an area adds capacity and resilience to that area.

Figure 8-5 Two RAPs per Cell with the Same Channel





Figure 8-6 Two RAPs per Cell on Different Channels

Multiple RAPs

If multiple RAPs are to be deployed, the purpose for deploying these RAPs needs to be considered. If the RAPs are being deployed to provide hardware diversity, the additional RAP(s) should be deployed on the same channel as the primary RAP to minimize the convergence time in a scenario where the mesh transfers from one RAP to another. When you plan RAP hardware diversity, consider the 32 MAPs per RAP limitation.

If additional RAPs are deployed to primarily provide additional capacity, then the additional RAPs should be deployed on a different channel than its neighboring RAP to minimize the interference on the backhaul channels.

Adding a second RAP on a different channel also reduces the collision domain through channel planning or through RAP cell splitting. Channel planning allocates different nonoverlapping channels to mesh nodes in the same collision domain to minimize the collision probability. RAP cell splitting is a simple, yet effective, way to reduce the collision domain. Instead of deploying one RAP with omnidirectional antennas in a mesh network, two or more RAPs with directional antennas can be deployed. These RAPs collocate with each other and operate on different frequency channels. This process divides a large collision domain into several smaller ones that operate independently.

If the mesh AP bridging features are being used with multiple RAPs, these RAPs should all be on the same subnet to ensure that a consistent subnet is provided for bridge clients.

If you build your mesh with multiple RAPs on different subnets, MAP convergence times increase if a MAP has to fail over to another RAP on a different subnet. One way to limit this process from happening is to use different BGNs for segments in your network that are separated by subnet boundaries.

Indoor Mesh Interoperability with Outdoor Mesh

Complete interoperability of indoor mesh APs with the outdoor ones is supported. It helps to bring coverage from outdoors to indoors. We recommend indoor mesh APs for indoor use only, and these APs should be deployed outdoors only under limited circumstances as described below.

**Caution**

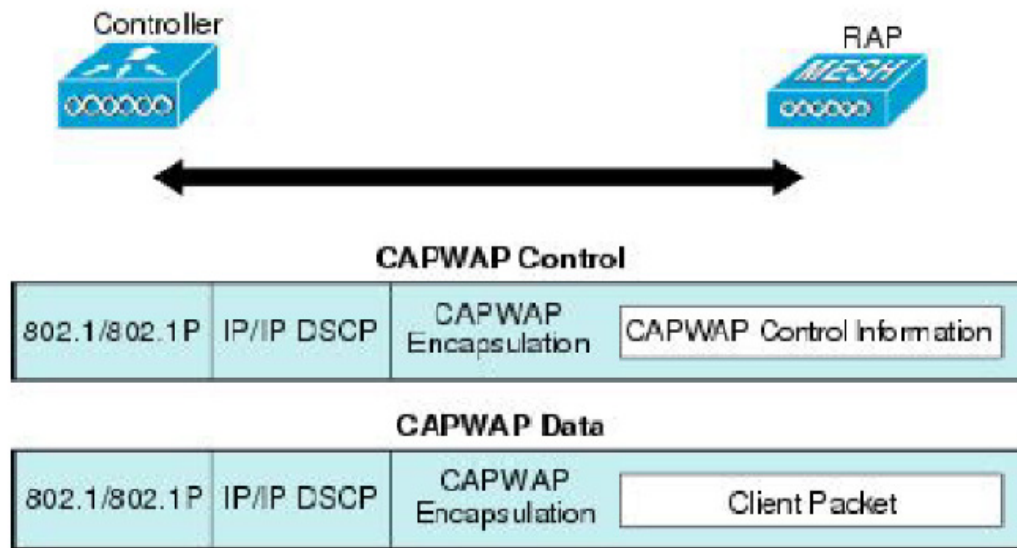
The indoor APs in a third-party outdoor enclosure can be deployed for limited outdoor deployments, such as a simple short haul extension from an indoor WLAN to a hop in a parking lot. The 1240, 1250, 1260, 2600, 3500e, and 3600 APs in an outdoor enclosure is recommended because of its robust environmental and temperature specifications. Additionally, the indoor APs have connectors to support articulated antennas when the AP is within an outdoor enclosure. Exercise caution with the SNR values as they may not scale and long-term fades may take away the links for these APs when compared to a more optimized outdoor 1500 series AP.

Mobility groups can be shared between outdoor mesh networks and indoor WLAN networks. It is also possible for a single controller to control indoor and outdoor mesh APs simultaneously. The same WLANs are broadcast out of both indoor and outdoor mesh APs.

Connecting the Cisco 1500 Series Mesh APs to the Network

This section describes how to connect the Cisco 1500 Series mesh APs to the network.

The wireless mesh terminates on two points on the wired network. The first location is where the RAP attaches to the wired network, and where all bridged traffic connects to the wired network. The second location is where the CAPWAP controller connects to the wired network; this location is where the WLAN client traffic from the mesh network connects to the wired network (see [Figure 8-7](#)). The WLAN client traffic from CAPWAP is tunneled at Layer 2, and matching WLANs should terminate on the same switch VLAN where the controllers are collocated. The security and network configuration for each of the WLANs on the mesh depend on the security capabilities of the network to which the controller is connected.

Figure 8-7 Mesh Network Traffic Termination**Note**

When an HSRP configuration is in operation on a mesh network, we recommend that the In-Out multicast mode be configured. For more details on multicast configuration, refer to the *Cisco Mesh Access Points, Design and Deployment Guide* at:

<http://www.cisco.com/en/US/docs/wireless/technology/mesh/7.3/design/guide/Mesh.html>.

Adding Mesh APs to the Mesh Network

This section assumes that the controller is already active in the network and is operating in Layer 3 mode.

**Note**

Controller ports that the mesh APs connect to should be untagged.

Before adding a mesh AP to a network, do the following:

- Step 1** Add the MAC address of the mesh AP to the controller's MAC filter. See the Adding MAC Addresses of Mesh Access Points to MAC Filter section.
- Step 2** Define the role (RAP or MAP) for the mesh AP. See the Defining Mesh Access Point Role section.
- Step 3** Verify that Layer 3 is configured on the controller.
- Step 4** Configure a primary, secondary, and tertiary controller for each mesh AP. Configure a backup controller. See the Configuring Backup Controllers section.
- Step 5** Configure external authentication of MAC addresses using an external RADIUS server. See the Configuring External Authentication and Authorization Using a RADIUS Server.
- Step 6** Configure global mesh parameters.
- Step 7** Configure universal client access.
- Step 8** Configure local mesh parameters.
- Step 9** Configure antenna parameters.

- Step 10** Configure channels for serial backhaul. This step is applicable only to serial backhaul APs.
 - Step 11** Configure the DCA channels for the mesh APs.
 - Step 12** Configure mobility groups (if desired) and assign controllers.
 - Step 13** Configure Ethernet bridging (if desired).
 - Step 14** Configure advanced features such as Ethernet VLAN tagging network, video, and voice.
-



VoWLAN Design Recommendations

This chapter provides additional design considerations for deploying voice over WLAN (VoWLAN) solutions. WLAN configuration specifics can vary depending on the VoWLAN device being used and the WLAN design. This chapter provides details about key RF and site survey considerations that are generally applicable to VoWLAN deployments, which were introduced in [Chapter 3, “WLAN RF Design Considerations.”](#)

Softphone applications are key VoWLAN solutions and they are available on a number of hardware and operating systems platforms. The Cisco Jabber™, application lets you access presence, instant messaging (IM), audio, video, voice messaging, desktop sharing, and conferencing. Jabber downloads for smartphones, tablets, and laptops along with information on design guides for each of the variants of Jabber can be found at: <http://www.cisco.com/web/products/voice/jabber.html>.

Antenna Considerations

The more demanding network requirements of VoWLAN impacts WLAN planning at all levels, including the choice of antenna. Key antenna considerations are as follows:

- Access point (AP) antenna selection
- Antenna placement
- Handset antenna characteristics

AP Antenna Selection

Cisco recommends a ceiling-mounted antenna solution for VoWLAN applications. Ceiling mounted antennas and APs with internal antennas are quick and easy to install. More importantly, they place the radiating portion of the antenna in open space, which allows the most efficient signal propagation and reception. Cisco APs with internal multiple antennas offer the easiest installation solution, plus the internal antennas provide a downward signal propagation pattern that is well suited for the majority of installations. The internal antenna solution is particularly well suited to the open spaces of enterprise environments.

Cisco offers a variety of multiple-input and multiple-output (MIMO) dual band, multiple element omni and patch antennas. These multiple element antennas are designed to take advantage of the Cisco AP technologies of Maximum Ratio Combining (MRC) and ClientLink. These technologies combine client phone packets, (as they are captured on the multiple antennas of the APs) into a single, combined signal that is stronger. The combined signal provides a better signal-to-noise ratio (SNR) between the phone’s transmitted packet and the general 2.4 or 5 GHz band noise. An important feature of MRC is that it

reduces the upstream packet error rate. Cisco APs use the multiple antennas and 802.11 ClientLink logic to deliver a higher energy packet to the client phone, which reduces the downstream packet error rate. These two features improve the mean opinion score (MOS) value of individual VoWLAN calls and the overall capacity of the Wi-Fi channel of the APs.

Cisco recommends that all antennas be placed 1 to 2 wavelengths away from highly reflective surfaces such as metal. The length of the 2.4 GHz waves is 4.92 inches (12.5 cm), and the length of the 5 GHz waves is 2.36 inches (6 cm). The separation of one or more wavelengths between the antenna and reflective surfaces allows the AP radio a better opportunity to receive a transmission and reduces the creation of nulls when the radio transmits. Orthogonal frequency-division multiplexing (OFDM), used by the 802.11g/n and 802.11a/n/ac specifications, helps to mitigate problems with reflections, nulls, and multipath. However, good antenna placement and the use of the appropriate antenna types provide a superior solution. The ceiling tile itself is a good absorber of signals transmitted into the area above the ceiling and reflected back into the coverage area.

For information on MRC see the IEEE report:

<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?reload=true&arnumber=1406225>

For information on ClientLink see the IEEE report on beamforming:

http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4558648&tag=1

Antennas come in a variety of types and form factors; no single type is best for all applications and locations. For additional information on the performance and part numbers of various antenna types, see the *Cisco Aironet Antennas and Accessories Reference Guide* at:

http://www.cisco.com/en/US/products/hw/wireless/ps469/products_data_sheet09186a008008883b.html

Cisco recommends using the Cisco Aironet dipole dual band AIR-ANT2524D series antenna when attaching dipole antennas to an AP with dual (2.4 GHz and 5 GHz) band support from the same external antenna port.

The Aironet dipole dual band antennas provide the advantages of:

- Support for simultaneous 2.4 and 5 GHz dual band transmission and reception (the same as the dual band omni and patch antennas). The gain on the Aironet dipole dual band antennas is 2.2 dBi for the 2.4 GHz band and 4 dBi for the 5 GHz band.
- Being small and coming in neutral colors of black, grey and white.
- Having an articulating and rotating base.

Antenna Orientation

Cisco recommends that, **for APs with multiple antennas, all the antennas be oriented in the same direction.**



Note

While APs are often depicted in marketing material showing the antennas arranged in multiple directions, as shown in [Figure 9-1](#), Cisco does not recommended this practice.

Figure 9-1 *AP With Antennas Arranged (incorrectly) in Different Directions*



The best MRC and ClientLink performance is obtained when all antennas of an AP are arranged in the same orientation, as shown in [Figure 9-2](#).

Figure 9-2 *AP With Antennas Arranged (correctly) in Same Orientation*



Having all four antennas of the AP in a flat, straight out position increases the overall throughput of the coverage cell by 2 Mbps when using single spatial stream 802.11n smartphones.

General Recommendations

Cisco recommends for optimum Wi-Fi coverage cell bandwidth and client application performance (for dipole antenna types of all forms) that:

- Each AP antenna port be populated with an antenna
- Each port must have the same antenna model
- Each antenna has the same orientation
- All of the antennas connected to the AP should be within two wave lengths of each other

The APs and the protocols they operate with are designed around MRC and ClientLink. Use an antenna system that follows these recommendations to capitalize on that technology and your AP hardware investment.


Higher gain antennas spread the signal on the horizontal plane, which creates a larger cell that also picks up additional noise. This results in a lower SNR that increases the packet error ratio. SNR is defined by the following criteria:

- **Signal**—The radiated energy transmitted from one radio that can be received uninterrupted by another radio. For Wi-Fi this means that the transmitting radio is sending 802.11 protocol packets that the receiving radio is able to decode.
- **Noise**—Transmitted energy in the frequency range of the receiving radio that cannot be decoded by that radio.

The larger the difference in energy between the protocol packet and the background noise, the better the reception of the protocol packet and the lower the packet error rate and bit error rate. Coverage area design involves using channels to create the lowest possible packet error rate while maintaining a high audio call capacity.

Higher gain antennas can also reduce the number of calls on a Wi-Fi channel because of the increased coverage area. For audio, a ceiling-mounted antenna is preferred over a wall-mounted patch because the human head and body attenuate 5 dB of the signal (see [Figure 9-3](#)). Ceiling mounted antennas are better positioned to avoid more of this head and body attenuation than most wall-mounted antennas.

Figure 9-3 *Head and Hand Attenuation*



Antenna Positioning

Ceiling-mounted antennas typically have better signal paths to handheld phones. The recommended coverage cell size takes into consideration the signal loss because of the attenuation of human heads and other obstacles. It is important to understand that the gain of antennas is reciprocal; gain applies equally to reception and transmission. Antenna gain is not an increase in transmitted power because the radio produces the transmitted power. The antenna is only a passive device. Gain is derived by focusing the signal of the radio into a direction, plane, and beam width, much in the same way a flashlight reflector focuses the light emanating from its bulb.

For further information on WLAN RF planning, see [Chapter 3, “WLAN RF Design Considerations.”](#)

Handset Antennas

For phones that integrate the antenna inside the body of the phone, the way the user holds the phone in their hand can influence signal attenuation by 4 dB. In some cases, a phone held against the head with the hand covering the antenna can result in a signal drop of 9 dB. The general rule for indoor deployments is that every 9 dB of signal loss reduces the coverage area in half. [Figure 9-3](#) shows an example of the difference in radiating power from a handset when held to the head.

The typical smartphone and tablet computer have a Wi-Fi antenna system with negative dB gain. The typical smartphone antenna is -3 or -4 dBi. The typical laptop has a positive gain from 0 to 2 dBi. This difference in antenna gain reflects in a difference in coverage range between smartphones, tablets and laptops at the same AP. For a smartphone or tablet to obtain the best application performance, the AP channel coverage should be designed to the Wi-Fi capabilities of the smartphones or tablets themselves. To provide optimum link quality between the smartphone, tablet or laptop and the AP, the AP should operate with ClientLink enabled. ClientLink is enabled by default with the Cisco wireless LAN controller (WLC).

Channel Utilization

The 802.11, 802.11b, 802.11g and 802.11n protocol specifications use the same 2.4 GHz band and therefore they must be able to interoperate with each other. This interoperability introduces additional 802.11 protection protocol logic overhead that reduces channel throughput. Many sites already have devices using the Wi-Fi frequencies of the 2.4 GHz band, but there are a number of foreign devices that can use these same frequencies. These foreign devices include Bluetooth, cordless phones, video game controllers, surveillance cameras, and even microwave ovens. Because the 2.4 GHz band is so crowded, and coupled with constraints on its channel allocation, Cisco recommends using the 5 GHz Wi-Fi band for new VoWLAN deployments. The channels available in the 5 GHz band are generally not used at most sites (see [Figure 9-4](#)). It is important to note that use of the 5 GHz UNII-2 channels for VoWLAN traffic requires the absence of radar. Cisco therefore recommends that there should be additional testing at any new site to determine whether a particular UNII-2 channel should be configured to be blocked. The reason is that if an AP detects radar on a channel during normal use, it must leave that channel within 200 ms.

Figure 9-4 Channel Utilization for 2.4 GHz Reporting



Before the installation of a Cisco Unified Wireless Network, the site can be tested for channel interference and utilization with tools from AirMagnet, Wild Packets, Cognio, and others. To aid in the design process, the *AP On-Demand Statistics Display* report generated by the Cisco Prime Infrastructure provides a spectrum review of:

- Client count versus RSSI

- Client count versus SNR
- Channel utilization

The ALOHAnet protocol defines a radio channel as full when channel utilization reaches 33 percent. This means the channel is busy enough that packets must wait for an open time slot before they are transmitted. The 46 percent channel utilization, as shown in [Figure 9-4](#), is above the channel utilization wireless packetized Aloha standard.

To reduce channel utilization in the 2.4 GHz band, Cisco recommends moving clients to 5 GHz and removing the legacy 1 Mbps and 2 Mbps data rates from the 2.4 GHz configuration when legacy devices are not part of the client makeup.

Dynamic Frequency Selection and 802.11h Requirements of the APs

The Federal Communications Commission (FCC) of the United States, the European Telecommunications Standards Institute (ETSI), and other regulatory agencies have their own requirements regarding the use of radio frequencies. Portions of the 5 GHz band have been and are currently being used for such things as weather radars. Although most 5 GHz radar systems generally use higher frequencies with shorter wavelengths, there are still systems in place that overlap with some Wi-Fi frequencies in the 5 GHz UNII-2 bands. In 2006, the FCC opened the frequencies in the 5.470 to 5.725 MHz range to unlicensed use. With these additional frequencies came a requirement to maintain an *interference-free* AP configuration. The AP must constantly monitor for radar pulses (typically from military, satellite, or weather stations) and use dynamic frequency selection (DFS) to automatically switch to a *clean* channel if radar is detected.

When radar is detected, the system must carry out the following:

- Stop packet transmission within 200 ms
- Stop control transmissions within 10 seconds
- Avoid transmission on the channel for 30 minutes
- Scan a new channel for 60 seconds before transmission

Because the radar avoidance requirements in the UNII-2 band can impact audio call quality, you should conduct a test for radar before going live with audio applications. Cisco Spectrum Expert is an excellent tool to test for the presence of radar on certain channels. If radar is detected during a Spectrum Expert test, the APs can then be configured to block use of those channels. For more information on Spectrum Expert, see: <http://www.cisco.com/en/US/products/ps9393/index.html>

5 GHz Band Channels

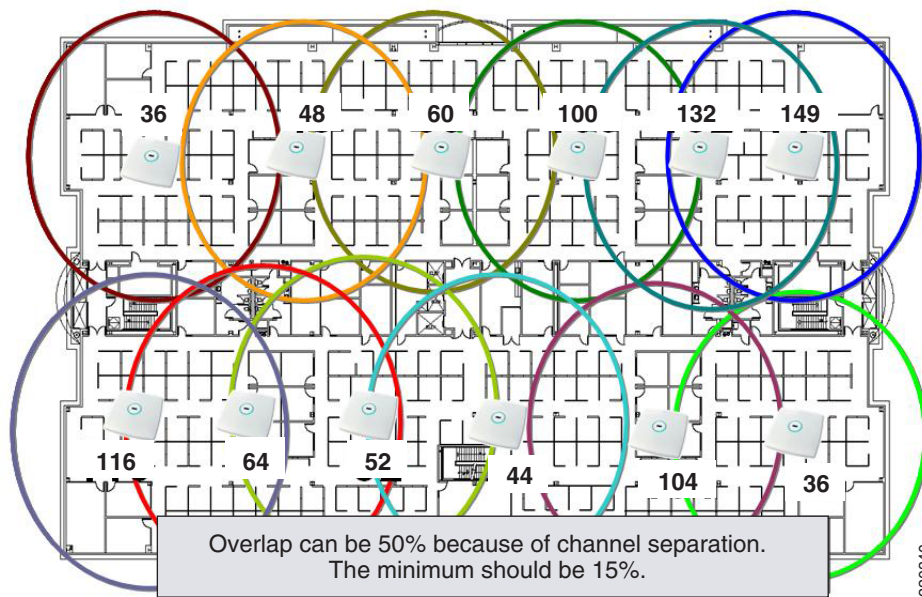
The DFS requirement includes the four original UNII-2 channels (52–64) and the new eight channels (100–116 and 132–140), while the 5 GHz band now has 20 channels. All of these channels are non-overlapping channels so all can be co-located. 2.4 GHz has only three non-overlapping channels. A design allowing co-located channels in a coverage area aggregates the number calls obtainable in a coverage area.




Note

See the Cisco website for current compliance information and also check with your local regulatory authority to find out what is permitted within your country.

A channel-based design can be implemented horizontally on a single floor, as shown in [Figure 9-5](#).


Figure 9-5 *Single Floor Channel Design*

In a multi-floor design, the channels can be separated vertically between floors to reduce the possibility of co-channel interference, as shown in [Figure 9-6](#).

Figure 9-6 Vertical Channel Separation

Call Capacity

The number of calls on a Wi-Fi channel is limited by a number of factors. First, the RF spectrum used by the AP and VoWLAN clients cannot be shielded from electromagnetic interference as shielded twisted-pair CAT 5 cable can. The closest Wi-Fi comes to segmentation is channel separation. The open, shared RF spectrum of 802.11 creates the possibility for high packet loss. Most of the packet loss is addressed through retransmission of 802.11 frames, which in turn causes jitter. [Figure 9-7](#) illustrates the packet loss relationship as a mean opinion score (MOS).

Figure 9-7 Effective Packet Loss Graphic

In the 802.11a specification as well as in 802.11g, the highest coverage range is achieved by the lowest data rate, which is 6 Mbps. For any given power level the lowest packet error rate is also 6 Mbps.

An acceptable coverage area for audio is an area that maintains a packet error rate of 5 percent or less. The MOS scores are ranked as follows:

- 4.4—Highest MOS score
- 4.3–4.0—*Very satisfied to Satisfied*
- 4.0–3.6—*Some users satisfied*

Figure 9-7 above shows that a packet error rate of 5 percent reduces the MOS to a level of *Some users satisfied* quality of speech.

The coverage area edge for a phone is the point in the coverage area that lowers the MOS rating to the *Very satisfied* category. This coverage area edge is referred to as a *cell edge* in this design guide. A cell edge with a 1 percent packet error rate is needed for audio because of the likelihood of multiple phone clients, data clients, co-channel interference, and other un-accounted for interferers. Cell edge and coverage design are defined in detail in other sections of this chapter.

If 802.11 and 802.11b are not required to support legacy 2.4 GHz Wi-Fi clients, Cisco recommends disabling the data rates of 1, 2, 5.5, and 11 MHz.

If these rates are disabled, one or more 802.11g data rates must be set to *required*. Cisco recommends that the 6 MHz data rate be set to required, but this depends on the cell size design requirements, which might require using a higher bit rate. If possible, an 802.11g-only network is recommended rather than a combined 802.11b/g network. Most data clients and phone clients recognize the data rates advertised by the AP in its beacons and probe response. Therefore, clients send their management, control, multicast, and broadcast packets at the required data rates as advertised by the AP, while they can send their unicast packets at any of the data rates advertised by the AP. Generally, unicast packets are sent at a data rate that provides the highest reliable rate for the link between the AP and client. Cisco APs are capable of sending unicast packets at a data rate that is unique for each ClientLink.

SNR is an important consideration for packet reception. The receiving radio is either the AP radio or the phone radio. The SNR is not likely to be the same at both radios of the link. SNR and multipath interference must be considered at the AP and at the cell edge. Path loss can be assumed to be the same at both ends of the link.

Cisco recommends that for audio applications the cell edge be determined by using the actual phone at the desired data rate. The audio packets sent between the AP and the phone in Wi-Fi applications are generally unicast real-time transport protocol (RTP) G.711 packets with a typical size of 236 bytes. The RTP packet is based on UDP and IP protocols, and therefore RTP is connectionless. The signal strength, SNR, data rate, and error rates of the phone call can be seen from the AP statistics, either on the autonomous AP or the controller-based CAPWAP AP.

Cisco also recommends that coverage testing be done with active calls. The two-way call provides the downstream (AP to client) packet size and the unicast packet type for ClientLink. The upstream (client to AP) provides the packets size and unicast packet type for MRC processing on the AP. When doing the client cell edge range testing, Cisco recommends testing a combination of smartphone, tablet and laptop models to the same AP from the same location and that the same square feet of space be used for all clients. This then means that the phones are not tested simultaneously because they could not all share the same space.

Figure 9-8 shows a sample of the client cell edge dBm values of a phone for 2.4 GHz and 5 GHz.

Figure 9-8 Client Edge RSSI -67 dBm with an SNR of 59 dB

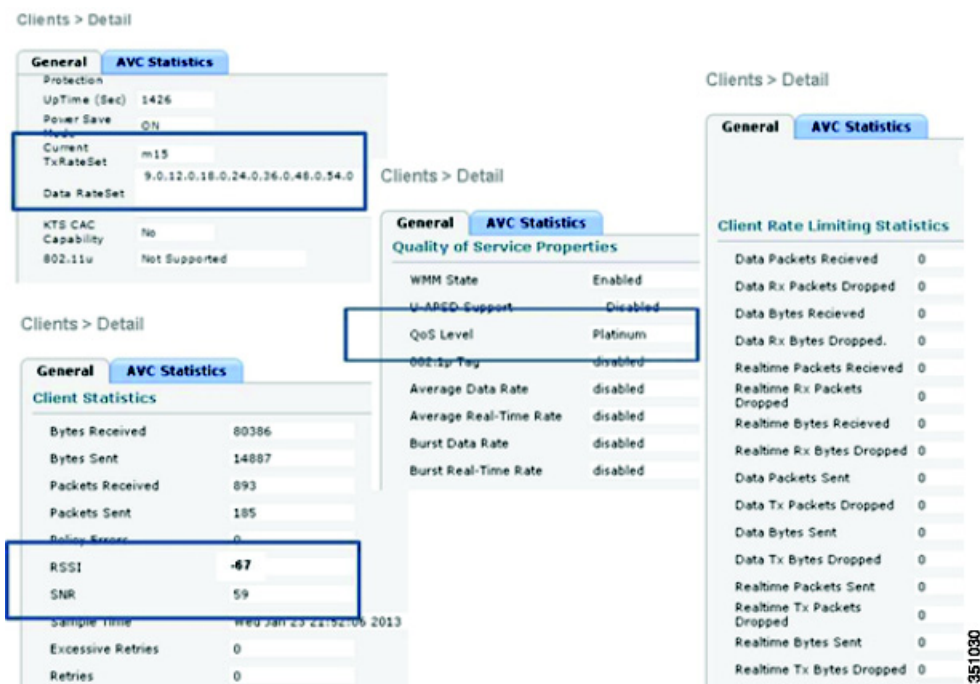


Figure 9-9 shows a decoded audio G.711 RTP packet. The packet, which originated on a Cisco 7960 desk phone, is downstream from the AP to a VoWLAN end-point. The over-the-air QoS marking is changed from the QoS baseline marking of 5 to a user priority of 6, which follows the 802.11e specification. Call statistics on the Cisco phone can be viewed on the phone or by browsing into the phone using the IP address of the phone. After that the cell edge dBm value can then be the benchmark value for tools that are better suited for surveying. A automated survey tool will expedite the coverage design of the site.

Figure 9-9 Sample VoWLAN Capture

When multipath interference is present at a location where signal level measurements are being taken, it is quite likely that the reported values will fluctuate from packet to packet. A packet can be as much as 5 dB higher or lower than the previous packet. It may take several minutes to obtain an average value for a given measurement location.

AP Call Capacity

A key part of the planning process for a VoWLAN deployment is to plan the number of simultaneous audio streams per AP.



Note

A call between two phones associated to the same AP is considered as two active audio streams.

When planning the audio stream capacity of the AP, consider the following points:

- The utilization of an unlicensed (shared) 802.11 channel is the real determinant for the number of simultaneous audio streams an AP can carry.
- Because the channel utilization and AP performance determine the number of audio streams, same channel and next channel separation are very important. Two APs in the same location, operating on the same channel, do not provide twice the number of audio streams. In fact, there can be fewer audio streams than a single AP would provide.
- Cell capacity or bandwidth determines the number of audio streams that can be simultaneously conducted.
- The QoS features supported in the handsets and VoWLAN deployment should be considered.

- Various handsets have different WLAN QoS features and capabilities that impact the features that are enabled in the WLAN deployment, and ultimately determine the per-AP audio call capacity of the AP. Most VoWLAN handsets provide guidance on the number of calls per AP supported by that phone; this should be considered a best-case figure for situations where the handset is able to use its optimal QoS features and has full access to the channel capacity.

The actual number of audio streams a channel can support is highly dependent on a number of issues, including environmental factors and client compliance to Wi-Fi Multimedia (WMM).

The table in [Figure 9-10](#) lists how Cisco Compatible Extensions benefit VoWLAN call quality.

Figure 9-10 Cisco Compatible Extension VoWLAN Features

How Cisco Compatible Extensions Benefits VoWLAN Call Quality	
Feature	Benefit
CCKM Support for EAP-Types	Locally Cached Credentials Means Faster Roams
Unscheduled Automatic Power Save Delivery (U-APSD)	More Channel Capacity and Better Battery Life
TSPEC-Based Call Admission Control (CAC)	Managed Call Capacity for Roaming and Emergency Calls
Voice Metrics	Better and More Informed Troubleshooting
Neighbor List	Reduced Client Channel Scanning
Load Balancing	Calls Balanced Between APs
Dynamic Transmit Power Control (DTPC)	Clients Learn a Power to Transmit At
Assisted Roaming	Faster Layer 2 Roams

220352

From [Figure 9-10](#) it can be seen that:

- Cisco Centralized Key Management (CCKM) provides faster client roaming for Extensible Authentication Protocol (EAP)-authenticated clients, which benefits audio call quality.
- Call Admission Control (CAC) also benefits audio call quality and can create bandwidth reservation for E911 and roaming calls.
- Assisted Roaming and Neighbor List benefit audio call quality and battery life.
- Voice metrics can benefit management.
- Unscheduled automatic power save delivery (U-APSD) and dynamic transmit power control (DTPC) benefit battery life.
- Load balancing and DTPC benefit audio call quality.

The Cisco Compatible Extensions Program provides third-party verification of Cisco Aironet wireless infrastructure products and wireless client devices from third-party companies. Several of the Cisco Compatible Extensions features have more than one benefit.

The amount of buffer memory, CPU speed, and radio quality are key factors of the performance of an AP radio. QoS features prioritize the audio and data traffic in the channel. For a further discussion of QoS, see [Chapter 5, “Cisco Unified Wireless QoS.”](#)

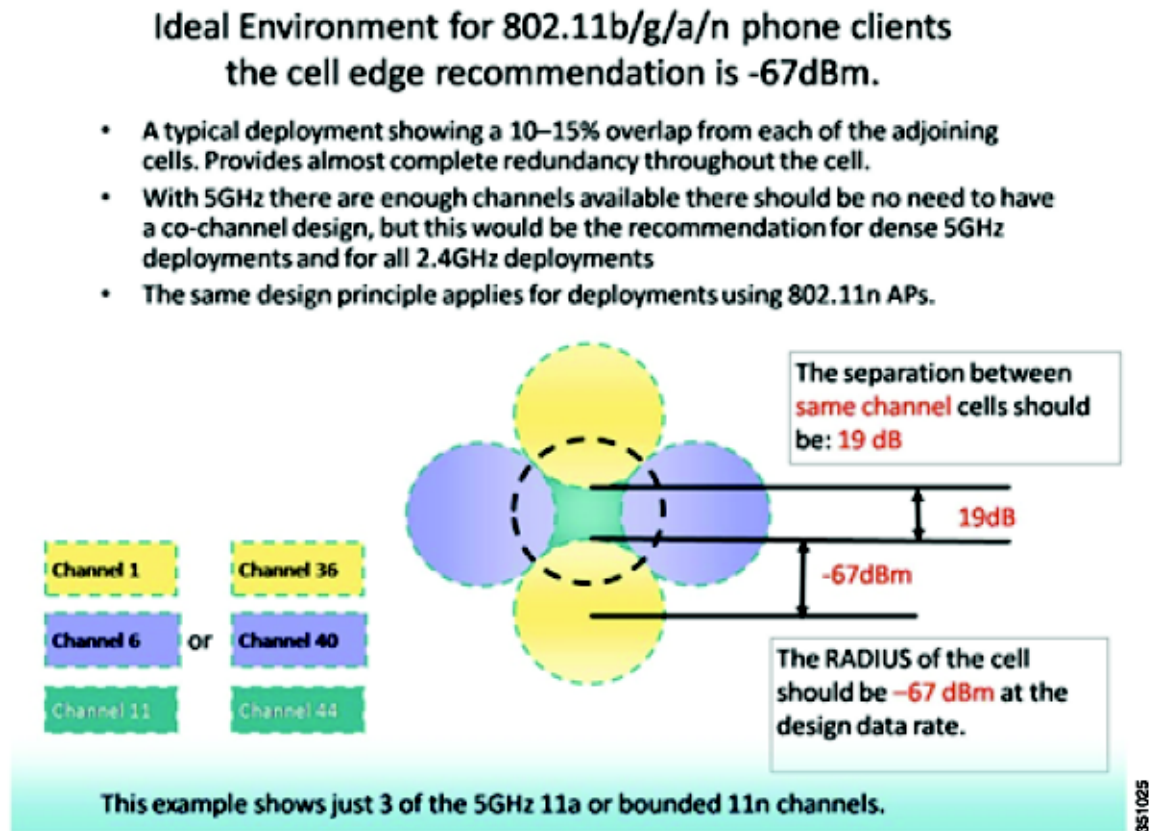
The 802.11e, WMM, and Cisco Compatible Extension specifications help balance and prevent the overloading of a cell with audio streams. CAC determines whether there is enough channel capacity to start a call; if not, the phone can scan for another channel. The primary benefit of U-ASPD is the preservation of WLAN client power by allowing the transmission of frames from the WLAN client to trigger the forwarding of client data frames that are being buffered at the AP for power saving purposes. The Neighbor List option provides the phone with a list that includes channel numbers and channel capacity of neighboring APs. This is done to improve audio call quality, provide faster roams, and improve battery life.

Cell Edge Design

Cisco guidelines for deploying 802.11b/g/a VoWLAN handsets recommend a design where a minimum power of -67 dBm is present at the cell boundary (see [Figure 9-11](#)). This practice creates cell sizes that are smaller than those used in data WLAN designs of the past. The -67 dBm threshold is a general recommendation for achieving a packet error of one percent, which requires an SNR value of 25 dB or greater (local noise conditions impact this requirement). Therefore, when determining the likely channel coverage area for a particular phone type, both signal strength and noise measured at the phone must be verified using the client statistics offered through the AP. See [Figure 9-10](#) for determining these values on the autonomous and CAPWAP APs.

The -67 dBm signal strength measurement has been used by 802.11b phone vendors for a number of years, and tests indicate that this same general rule of measurement also works well for 802.11g/n and 802.11a/n phone clients.

Figure 9-11 Cell Edge Measurements

**Note**

The -86 dBm separations shown in Figure 9-11 are simplified and are considered ideal. It is not likely that this 19 dBm of separation can be achieved in most deployments. The most important RF design criteria are the -67 dBm cell radius and the 20 percent recommended overlap between cells. Designing to these constraints optimizes channel separation.

For 5 GHz cells, there is less concern about same channel separation because of the number of available non-overlapping channels. There are 20 channels in the 802.11a 5 GHz band so a two-channel separation is almost always possible. In contrast, the 2.4 GHz band has only three channels that do not overlap in frequency.

For both the 5 GHz and 2.4 GHz bands the cell edge must be at the floor level where a packet error rate of 1 percent is maintained at the highest data rate desired for a given channel. The data rate is 72 Mbps for an 802.11n on 2.4 GHz with a one spatial stream client.

The data rate is 150 Mbps for an 802.11n client on the 5 GHz band with a 40 MHz wide channel and with a one spatial stream client. A laptop running a softphone application such as Jabber can support three spatial streams and have a data rate to 450 Mbps on a 5 GHz, 40 MHz wide channel. Both an 802.11a client and an 802.11n client that is only 20 MHz wide and supporting one spatial stream can share Wi-Fi channel access on a 40 MHz wide channel with an 802.11ac three spatial stream client on a 80 MHz wide channel.

This type of client mixture and protocol mixture is part of the 802.11 specification. The *compatibility* for this type of client mixture on the same Wi-Fi frequencies is part of the 802.11n and 802.11ac specifications.

The major design question is how to define the coverage area for bandwidth and call capacity. Audio call capacity is about the same for 802.11n and 802.11ac as it was for 802.11g and 802.11a. This is because of the packet size of the audio G.711 or G.722 frame, which with AES encryption is less than 300 bytes. The small packet size and the ACK logic of the 802.11 specification creates a large overhead compared to large streaming applications. A video call generates both small audio packets and large video packets. The video packets are highly compressed and therefore spaced out in comparison to the audio. Cisco recommends as a guideline to establish the cell edge of coverage. Measure the distance from the AP that the phone is when the RSSI value of phone on the AP is -67 dBm.

802.11g and 802.11a phone clients can be capable of rates up to 54 Mbps. Current chip sets support 54 Mbps, but transmit power capabilities differ. Cisco highly recommends that all links between phone clients and APs be established using matching transmit power levels (see [Dynamic Transmit Power Control](#), page 9-17).

Coverage cells can be created for specific data rates. For a high density deployment or a deployment where a large number of calls are required within a small floor space, 802.11a is recommended because of the number of channels and the 54 Mbps data rate. The lower data rates in 802.11a can be disabled, the 24 Mbps data rate can be set to *required*, while the 36 to 54 Mbps data rates can be left enabled.

After setting the cell edge of -67 dBm, determine where the error rate of 1 percent occurs, and then examine the SNR value.

The -67 dBm cell edge can be determined as follows:

- Set the phone to its desired transmit power.
- Set the AP to a matching transmit power.
- Place the AP and the desired antenna in the location where the phone will be used.
- With an active call, or while sending and receiving packets equal in size to the G711 codec, measure the signal level out to the -67 dBm cell edge.

Carefully examine the data sheets of the particular phone device to determine the transmit power levels and data rates supported by the phone device in a particular Wi-Fi band. The data sheets for Cisco Unified Wireless IP Phones can be found at:

<http://www.cisco.com/en/US/products/hw/phones/ps379/index.html>.

The 2.4 GHz maximum transmit power levels vary on different channels and with different AP models. The 5 GHz maximum transmit power levels vary by model. The Cisco Aironet AP data sheets should be carefully reviewed to determine which AP model supports which data rates. [Figure 9-12](#) shows an example of the maximum 5 GHz transmit power in dBm by channel.

Figure 9-12 Channel Power Assignment

UNII-1				UNII-2				UNII-3				802.11a
36	40	44	48	52	56	60	64	149	153	157	161	
14	14	14	14	17	17	17	17	17	17	17	17	17
Extended UNII-2												802.11b
100	104	108	112	116	120	124	128	132	136	140		
17	17	17	17	17	17	17	17	17	17	17		

The maximum permissible transmit power across the 5 GHz band varies by as much as 6 dB. This means that when using the maximum allowed transmit power throughout a site that allows all channels, there will not be equal cell coverage on all channels. It also means that if dynamic channel selection is used, the cell coverage edge can change based on the channel number. However, dynamic channel selection can be tuned. The default mode of dynamic channel selection accounts for the difference of maximum transmit power level by channel.

Cell transmit power on all APs should not exceed the maximum or desired transmit power of the phones. If the phone's maximum or set transmit power is 13 dBm, Cisco recommends that all APs have a maximum transmit power of 13 dBm. Therefore, the maximum transmit power on the AP should be set to an equal level or, if that is not possible, the next higher transmit power level. Equal transmit power is recommended to avoid one-way audio problems. The AP generally has better receiver sensitivity and diversity support than the phone, so it should be able to receive the slightly lower strength phone signal. See [Dynamic Transmit Power Control](#), page 9-17 for more information on equal transmit powers.

Dual Band Coverage Cells

[Chapter 3, "WLAN RF Design Considerations"](#) describes 2.4 GHz and 5 GHz channel coverage design. For a dual mode AP to provide equal cell coverage on both the 2.4 GHz and 5 GHz channels, the 2.4 GHz channel must have an equal (or usually lower) transmit power than the 5 GHz channel. At most sites the noise level in the SNR formula is lower by up to 10 dB. The receiver sensitivity of 802.11g radios is generally 2 dBm better than the same data rate on 802.11a radios. For example, the Cisco 7921G phone data sheet lists the receive sensitivity of -78 dBm at the data rate of 36 Mbps for 802.11g, and -76 dBm for 802.11a. Therefore, given the anticipated better noise floor of 10 dB, the 802.11a cell can do better by 8 dBm. Other details such as the difference in path loss between 802.11g and 802.11a keep this from being a direct ratio. However, if the same coverage cells are desired, reducing the 802.11g network by one or two power levels from the 802.11a network power levels should accomplish this goal.

Dynamic Transmit Power Control

By default, Cisco Aironet APs have dynamic transmit power control (DTPC) enabled. DTPC is automatic with Cisco WLCs but must be configured on autonomous APs.

The objective of DTPC is to reduce the chance of one-way audio because of an imbalance of transmit power between the AP and the Wi-Fi radio of the client. DTPC accomplishes this by:

- Setting the phones transmit power to match the transmit power of the APs
- Having APs advertise their transmit power for the clients to learn

DTPC allows phones to automatically adjust their transmit power to that of the APs. In the example shown in [Figure 9-13](#), this means that the phone changes its transmit power level from 5 mW to 100 mW.

Figure 9-13 Client and AP Power Matching



The licensing requirements of 802.11 do not require clients to have a minimum transmit power and few if any Wi-Fi devices use the maximum transmit powers allowed by regulations. With a typical Wi-Fi device, the maximum transmit power capability is at or below 100 mW. This is because Wi-Fi specifications do not require APs and clients to have matching power levels during associated connections between themselves. There will always be the possibility that for a short period of time,

while associated, they might not be in the coverage range of each other but still be associated. If this happens during an active call there is a loss of audio. If the transmit power levels are not equal during an active call then there is audio loss. Several 802.11 mechanisms help to maintain the connection between the AP and the phone, one being that they can negotiate a slower data rate. Slower data rates generally have higher transmit powers than higher data rates. The slower data rates should be avoided in dense deployments. This is because when a coverage cell needs high throughput and capacity, the slower data rates for the high packet count phone calls lowers the throughput for all clients on that Wi-Fi channel and AP.

Cisco highly recommends that the maximum configured transmit power on APs be no higher than the maximum transmit power the client phones support. Because the current Cisco APs support ClientLink, Cisco highly recommends that ClientLink be configured. ClientLink dynamically creates a directed signal towards selected clients. The ClientLink logic changes the signal prorogation on directed packets but not on broadcast or multicast packets. ClientLink removes the typically omni antenna horizontal signal prorogation with equal signal energy in all directions. Signal energy is increased in the direction for the selected clients. The directed signal increases the signal energy at the selected client, improving downstream signal quality at the phone. This improves the MOS value of the call. Improving the MOS value reduces retries and improves the throughput in the coverage area for all clients. Because this is a shaped signal that is directed to a specific location, there is reduced signal in the remaining coverage areas of the AP. This improves the performance of the channel in areas where there is channel overlap with broadcast and multicast packets with other APs.

Cisco recommends that each model of phone be tested for its Wi-Fi coverage range. The WLC reports the receive signal strength indicator (RSSI) of each client at the AP to which the phone is associated. The value shown in the RSSI field is the signal strength of a packet transmitted from the phone to the AP. The value indicates how strong the packet transmitted by the phone was at when it was received at the AP. It is recommended to check the coverage range of the phones and that the phones be placed at the estimated coverage edge of the AP. Then check the RSSI when a phone is on an active call. The goal is that at the cell edge (recommended -67dBm RSSI) the packets are sent at a high data rate. See [Figure 9-11](#) for a reference to the cell edge for VoWLAN Wi-Fi coverage area range. The value of -39 shown in the figure is a very strong signal that is seen when the client phone or device is within a few feet of the AP.

Testing the phone's coverage has become more important with the advent of smartphones and tablets. Because the Wi-Fi feature sets of these devices are typically for the consumer market, these devices typically have few 802.11 features that are considered to support enterprise. The consumer orientation of most smartphones and tablets does not support DTPC. Therefore, Cisco recommends that the maximum transmit power for 2.4 GHz and 5 GHz be a dBm value that matches the 2.4 GHz and 5 GHz band maximum transmit power of your weakest smartphone or tablet. This WLC field value limits the transmit signal power of the APs, thereby helping to maintain a balance in range of the phone to the AP.

802.11r and 802.11k Features

IEEE 802.11k and 802.11r are key industry standards that enable seamless Basic Service Set (BSS) transitions in the WLAN environment. With WLAN 7.2 release, Cisco supports the 802.11r secure authentication *Fast Transition* protocol. The IEEE 802.11k specification was ratified in June 2008. The IEEE 802.11r specification was ratified in July 2008. 802.11r specification follows the 802.11e security specification of June 2004.

For a brief description of the 802.11k specification see:

http://en.wikipedia.org/wiki/IEEE_802.11k-2008

For the 802.11k specification see:

<http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4544755>

For a brief description of the 802.11r specification see:

http://en.wikipedia.org/wiki/IEEE_802.11r-2008

For the IEEE 802.11r specification see:

<http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=04573292>

802.11k and 802.11k-enabled client devices send a request for a list of neighbor APs (a *Neighbor List*) from the APs they are currently associated with. The request is in the form of an 802.11 management frame known as an *action packet*. The AP responds with an action packet that contains a Neighbor List of APs on the same WLAN along with their Wi-Fi channel numbers.

From the response action packet the 802.11k client learns which APs are candidates for the next roam. The use of 802.11k radio resource management (RRM) algorithms allows smartphones to roam efficiently and quickly: a requirement for good call quality in an enterprise environment where on-call roams are common.

Cisco recommends that the 802.11k be configured in the WLC to enable radio resource management (RRM) to provide both 2.4 GHz and 5 GHz AP channel numbers in the Neighbor List response packets. Cisco also recommends the use of 5 GHz band Wi-Fi channels for not only VoWLAN calls but for all applications and devices.

With information from the Neighbor List, 802.11k clients do not need to probe all of the 2.4 GHz and 5 GHz channels to find an AP they can roam to. Not having to probe all of the channels reduces channel utilization on all channels, thereby increasing bandwidth on all channels. It also reduces roam times and improves the decisions made by the client. Additionally, it increases battery life of the device because it is neither changing the radio configuration for each channel nor sending probe requests on each channel. This prevents the devices from having to process all of the probe response frames.

The 802.11r and 802.11e specifications both support the same authentication types: EAP-FAST, LEAP, EAP-TLS, EAP-TTLS, EAP-SIM, and PEAP versions 1 and 2. This security feature allows an 802.11r-enabled client to authenticate securely to an AP in an exchange of only four packets. Two of the packets are sent over the Ethernet wires that connect the APs to each other. The other two packets are sent on the Wi-Fi channels of each AP. This allows the 802.11r client to be authenticated securely to the AP that it is going to roam to before it actually roams. The result is the 802.11r client can be sending and receiving data, video, and audio packets after the roam without the delay of the authentication process. Because the 802.11 header is changed by the addition of the 802.11r parameters, the WLAN for 802.11r clients cannot be shared with clients that are not 802.11r-aware. This means that all clients that have the SSID assigned by the WLAN with 802.11r enabled must have Wi-Fi radio firmware that is aware of the 802.11r element in association packets. Limitations to 802.11r fast roaming are:


- It is not supported on APs in autonomous mode
- Roaming between local authentication and central authentication WLAN is not supported

Cisco recommends that you use the 802.11r specification as it improves roam times because of a reduction in the number of packets sent over the Wi-Fi channel between a client that is already authenticated to the WLAN.

Interference Sources Local to the User


Interference can be local to the user but it is also likely to affect nearby users. Bluetooth is a popular RF protocol used in personal area networks that interferes with 2.4 GHz Wi-Fi channels. [Figure 9-14](#) shows that the actual Bluetooth signal does span all of the 2.4 GHz channels used by 802.11b/g clients. This graphic is taken from an 802.11g audio call with a Bluetooth headset linked to the phone. [Figure 9-15](#) also shows the jitter caused by the Bluetooth headset.

Figure 9-14 Signal Pattern in the 802.11b/g 2.4 GHz Spectrum of a Typical Bluetooth Earpiece



In [Figure 9-14](#) the **purple** line shows the Max Hold, the maximum transmit power that was reached during the test. The **yellow** line shows the maximum transmit power in the last sample period of ten seconds. The **turquoise** line shows the average transmit power over the period of the test. The vertical dashed **blue** lines separate the three non-overlapping 802.11b/g channels (Ch1, Ch6, and Ch11). The charting is from 2.400 GHz on the left to 2.500 GHz on the right. From the right edge of the Ch11 vertical blue line is the part of the 802.11 spectrum used in Europe and Japan. This capture was done with the AP and clients configured for the North American regulatory domain. This graph shows that the Bluetooth earpiece was easily transmitting outside of FCC regulations.

Notice that the Bluetooth signal is very narrow. Bluetooth transmits data on a single MHz of frequency, stops the transmission, moves to another frequency in the 802.11 2.4 GHz band, and then transmits data. This is repeated continually. The 802.11b and 802.11g signals are sent with a combined 22 MHz of frequency. The radio remains on that 22 MHz of frequency. This grouping of 22 MHz is referred to as the channel. The Max Hold line shows how strong the Bluetooth is while in search mode. The signal level is above that of a 50 mW (17 dBm) OFDM 802.11g radio. A signal of this strength and duration causes 802.11b/g phones to drop the VoWLAN call. Lesser strength Bluetooth signals cause jitter, resulting in a lower MOS value. [Figure 9-15](#) shows an example of an Ethereal jitter analysis of three simultaneous phone calls, each using a Bluetooth earpiece.

Figure 9-15 Jitter Analysis Example

All three calls were on the same AP and were to three other phones also on the same AP.


For information on interference with Wi-Fi and Bluetooth, see the IEEE report:

<http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6105779>

The factors that affect the impairments introduced to a Bluetooth TDM packet when colliding with Wi-Fi OFDM include:

- Relative power
- Bandwidth
- Mutual overlap
- Number of colliding OFDM signals

Simulations were performed on the effects of interference between a sample Wi-Fi OFDM packet and Bluetooth signals, as shown in Figure 9-16. The figure shows the *Normal* GMSK Bluetooth undistorted signal TDM characteristics. On the left is frequency versus time (MHz), and on the right is I/Q amplitudes.

Figure 9-16 IEEE Waveform Simulations

As shown above in the [Figure 9-16](#), the 625-second long hopping Bluetooth packet can interfere with more than one OFDM packet at a time, especially whenever high-rate OFDM mode packets (where the length of the packet is much shorter than that of Bluetooth) are subject to the collision.



Cisco Unified Wireless Network Guest Access Services

The introduction of wireless LAN (WLAN) technologies in the enterprise has changed the way corporations and small-to-medium businesses function by freeing staff and network resources from the constraints of fixed network connectivity.

WLAN has also changed how individuals access the Internet and their corporate networks from public locations. The advent of public WLAN hotspots has caused mobile workers to become accustomed to being able to access their corporate network from practically anywhere.

Introduction

The paradigm of public access has extended to the enterprise itself. Our highly mobile, information-on-demand culture requires on-demand network connectivity. For this reason, enterprise guest access services are becoming increasingly important and a necessity in the corporate environment.

While there is broad recognition that guest networking is becoming increasingly important, there is also well-founded apprehension over how to safeguard internal corporate information and infrastructure assets. When implemented correctly, an enterprise that implements a guest access solution will most likely improve their overall security posture as a result of the network audits associated with the implementation process.

In addition to overall improved security, implementing a guest access network offers these additional general benefits.

- Authentication and authorization control of guests based on variables including date, duration, and bandwidth
- An audit mechanism to track who is currently using, or has used, the network

Additional benefits of a wireless-based guest access include the following:

- It provides wider coverage by including areas such as lobbies and other common areas that otherwise might not have been wired for network connectivity.
- It removes the need for designated guest access areas or rooms.

Scope

Several architectures can be implemented to offer guest access in the enterprise. It is not the goal of this chapter to cover all possible solutions. Instead, this chapter focuses on the implementation of wireless guest networking using the Cisco Unified Wireless Network solution. For more information on deploying wired and wireless Guest Access services in other topology scenarios, see:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Network_Virtualization/GuestAcc.html.

Wireless Guest Access Overview

Ideally, the implementation of a wireless guest network uses as much of an enterprise's existing wireless and wired infrastructure as possible to avoid the cost and complexity of building a physical overlay network. Assuming this is the case, the following additional elements and functions are needed:

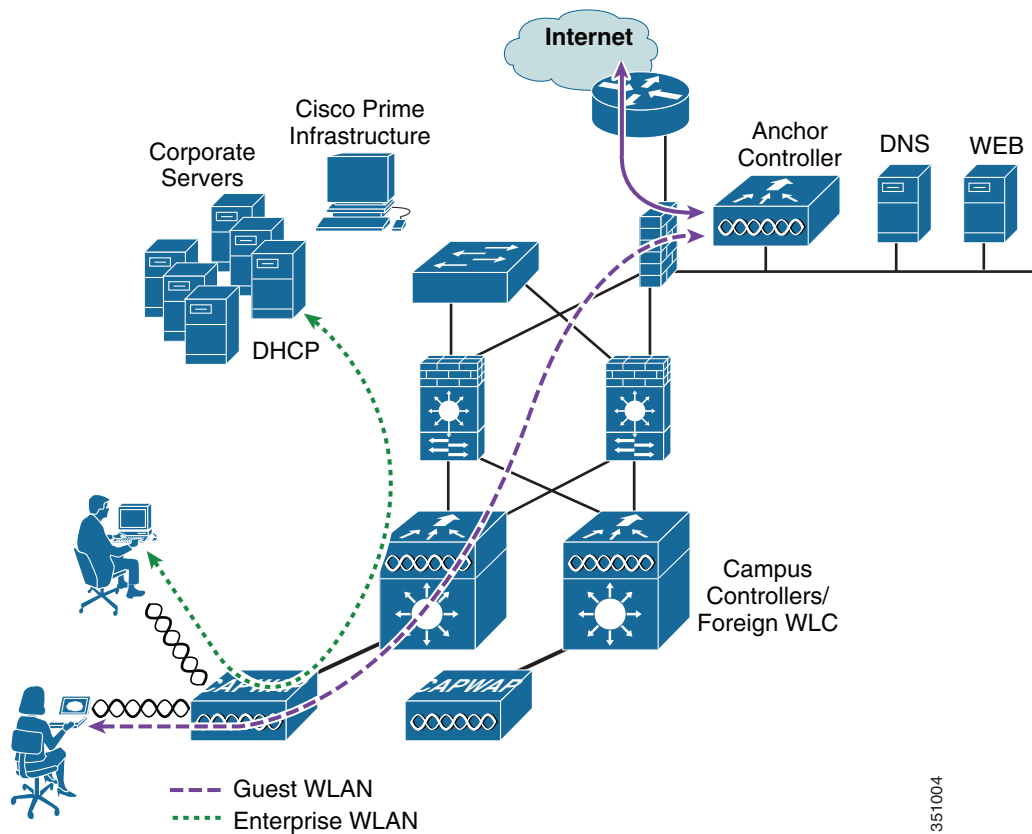
- A dedicated guest WLAN/SSID—Implemented throughout the campus wireless network wherever guest access is required.
- Guest traffic segregation—Requires implementing Layer 2 or Layer 3 techniques across the campus network to restrict where guests are allowed to go.
- Access control—Involves using imbedded access control functionality within the campus network or implementing an external platform to control guest access to the Internet from the enterprise network.
- Guest user credential management—A process by which a sponsor or lobby administrator can create temporary credentials in behalf of a guest. This function might be resident within an access control platform or it might be a component of AAA or some other management system.

Guest Access using the Cisco Unified Wireless Network Solution

The Cisco Unified WLAN solution offers a flexible, easy-to-implement method for deploying wireless guest access by using Ethernet in IP (RFC3378) within the centralized architecture. Ethernet in IP is used to create a tunnel across a Layer 3 topology between two WLC endpoints. The benefit of this approach is that there are no additional protocols or segmentation techniques that must be implemented to isolate guest traffic from the enterprise.

See [Figure 10-1](#) for an example of guest access topology using a centralized WLAN architecture

Figure 10-1 Centralized Controller Guest Access



As illustrated in [Figure 10-1](#) the anchor controller is located in the enterprise DMZ where it performs an “anchor” function. The anchor controller is responsible for terminating EoIP tunnels that originate from other campus controller throughout the network. These “foreign” controllers are responsible for termination, management, and standard operation of the various WLANs provisioned throughout the enterprise, including one or more guest WLANs. Guest WLANs are transported via an EoIP tunnel to the anchor controller. Specifically, guest WLAN data frames are encapsulated using CAPWAP from the AP to the foreign controller and then encapsulated in EoIP from the foreign management system to a guest VLAN defined on the anchor WLC. In this way, guest user traffic is forwarded to the Internet transparently, with no visibility by, or interaction with, other traffic in the enterprise.

WLAN Controller Guest Access

The Guest Access solution is self-contained and does not require any external platforms to perform access control, web portal, or AAA services. All these functions are configured and run within the anchor controller. However, the option exists to implement one or all of these functions externally and is discussed later in the chapter.

Supported Platforms

The anchor function, which includes tunnel termination, web authentication, and access control is supported on the following WLC platforms (using version 6.0 or later):

- WLC 5508
- WiSM-2
- WLC 7500

The following WLC platforms *cannot* be used for anchor functions, but can be used for standard controller deployments and guest mobility tunnel origination (foreign WLC) to a designated anchor controller(s):

- Cisco WLAN Controller Module for Integrated Service Routers (ISR-SM)
- Cisco 2504

Auto Anchor Mobility to Support Wireless Guest Access

Auto anchor mobility, or guest WLAN mobility, is a key feature of the Cisco Unified Wireless Network solution. It offers the ability to map a provisioned guest WLAN to one or more (anchor) WLCs by using an EoIP tunnel. Auto anchor mobility allows a guest WLAN and all associated guest traffic to be transported transparently across an enterprise network to an anchor controller that resides in the Internet DMZ (see [Figure 10-2](#)).

Figure 10-2 Auto Anchor EoIP Tunnels

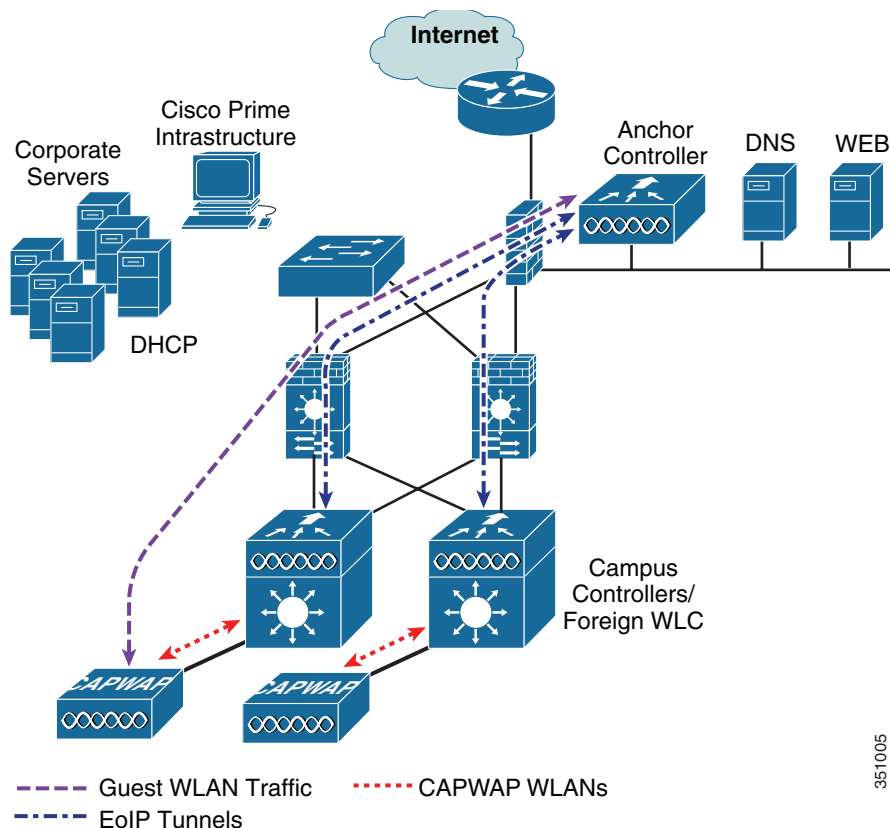
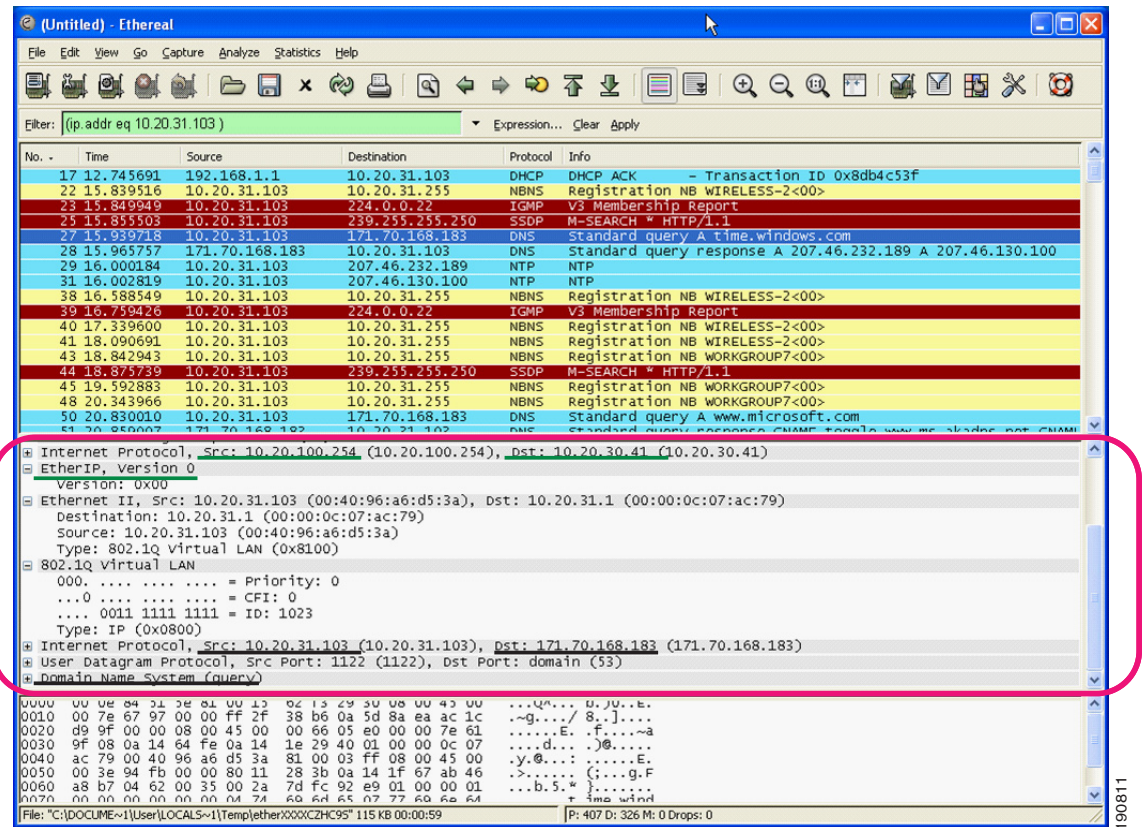


Figure 10-3 shows a sniffer trace of an Ethernet in IP tunnel (highlighted) between a foreign controller with a guest WLAN provisioned and an anchor controller that is performing local web authentication. The first IP detail shown represents the Ethernet in IP tunnel between the foreign and anchor controllers. The second IP detail is that of guest traffic (in this case, a DNS query).

Figure 10-3 Sample Ethernet in IP Sniffer Trace



Anchor Controller Deployment Guidelines

This section provides guidelines for deploying an anchor controller to support wireless guest access.

Anchor Controller Positioning

Because the anchor controller is responsible for termination of guest WLAN traffic and subsequent access to the Internet, it is typically positioned in the enterprise Internet DMZ. In doing so, rules can be established within the firewall to precisely manage communications between authorized controllers throughout the enterprise and the anchor controller. Such rules might include filtering on source or destination controller addresses, UDP port 16666 for inter-WLC communication, and IP protocol ID 97 Ethernet in IP for client traffic. Other rules that might be needed include the following:

- TCP 161 and 162 for SNMP
- UDP 69 for TFTP
- TCP 80 or 443 for HTTP, or HTTPS for GUI access
- TCP 23 or 22 for Telnet, or SSH for CLI access

Depending on the topology, the firewall can be used to protect the anchor controller from outside threats.

For the best possible performance and because of its suggested positioning in the network, it is strongly recommended that the guest anchor controller be dedicated to supporting guest access functions only. In other words, the anchor controller should not be used to support guest access in addition to controlling and managing other CAPWAP APs in the enterprise.

DHCP Services

As previously described, guest traffic is transported at Layer 2 via EoIP. Therefore, the first point at which DHCP services can be implemented is either locally on the anchor controller or the controller can relay client DHCP requests to an external server. See [Guest Access Configuration, page 10-12](#) for configuration examples.

Routing

Guest traffic egress occurs at the anchor controller. Guest WLANs are mapped to a dynamic interface/VLAN on the anchor. Depending on the topology, this interface might connect to an interface on a firewall, or directly to an Internet border router. Therefore, a client's default gateway IP is either that of the firewall or the address of a VLAN/interface on the first hop router. For ingress routing, it is assumed the guest VLAN is directly connected to a DMZ interface on a firewall or to an interface on a border router. In either case, the guest (VLAN) subnet is known as a directly connected network and advertised accordingly.

Anchor Controller Sizing and Scaling

The most cost-effective platform to support guest networking, in most enterprise deployments, is the Cisco 5508 Series controller. Assuming the controller is being deployed to support guest access with EoIP tunnel termination only, the 5508 with support for 12 APs is sufficient because it is assumed the controller is not going to be used to manage APs in the network.

A single 5508 Series controller can support EoIP tunnels from up to 71 foreign controllers within the enterprise. Additionally, the 5508 controller supports up to 7,000 simultaneous users and has a forwarding capacity of 8 Gbps.

The selection of the guest anchor controller is a function of the amount of guest traffic, as defined by the number of active guest client sessions, or as defined by the uplink interface capacity on the controller, or both.

Total throughput and client limitations per guest anchor controller are as follows:

- Cisco 2504 Wireless LAN Controller – 4 * 1 Gbps interfaces and 1000 guest clients
- Cisco 5508 Wireless LAN Controller (WLC) – 8 Gbps and 7,000 guest clients
- Cisco Catalyst 6500 Series Wireless Services Module (WiSM-2) – 20G bps and 15,000 clients
- Cisco 7500 Wireless LAN Controller (WLC) – 10 Gbps and 20,000 clients

Anchor Controller Redundancy

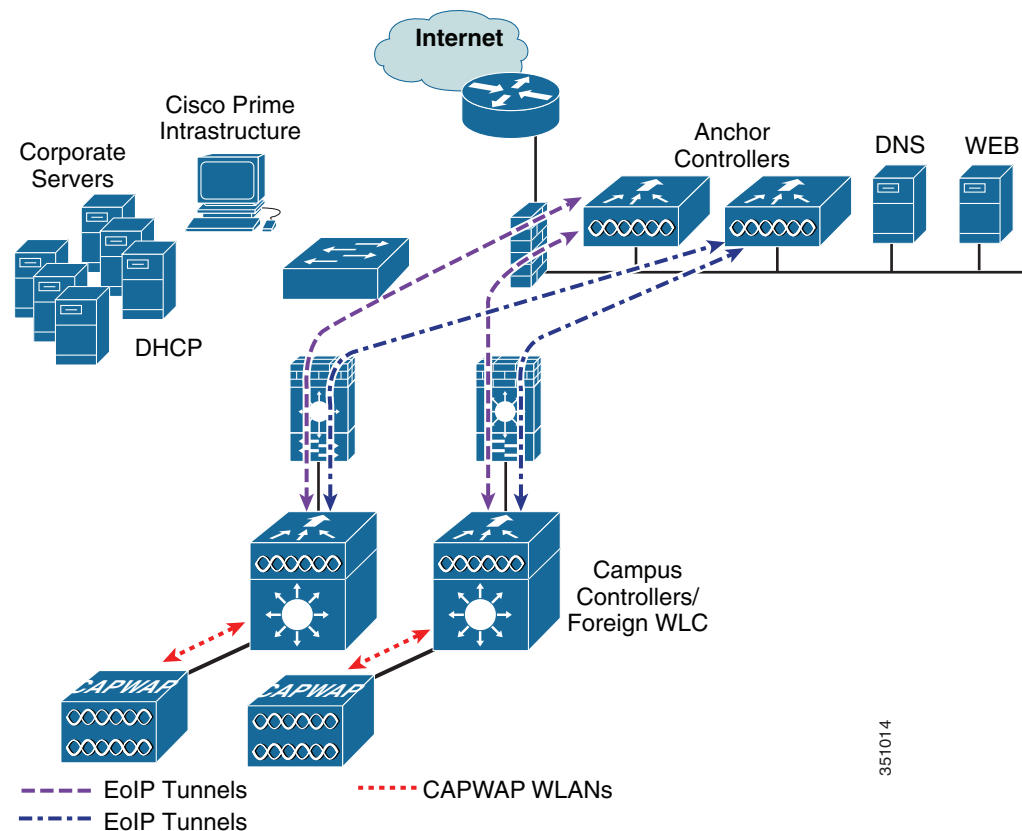
Beginning with Release 4.1 of Cisco Unified Wireless Network solution software, a “guest N+1” redundancy capability was added to the auto anchor/mobility functionality. This feature introduced an automatic ping function that enables a foreign controller to proactively ping anchor controllers to verify control and data path connectivity. In the event of failure or an active anchor becomes unreachable, the foreign controller does the following:

- Automatically detects that the anchor has become unreachable
- Automatically disassociates any wireless clients that were previously associated with the unreachable anchor
- Automatically re-associates wireless client(s) to an alternate anchor WLC

With guest N+1 redundancy, two or more anchor WLCs can be defined for a given guest WLAN.

Figure 10-4 shows a generic guest access topology with anchor controller redundancy.

Figure 10-4 Guest Access Topology with Guest Anchor N+1 Redundancy



Keep in mind the following in regards to guest N+1 redundancy:

- A given foreign controller load balances wireless client connections across the list of anchor controllers configured for the guest WLAN. There is currently no method to designate one anchor as primary with one or more secondary anchors.
- Wireless clients that are associated with an anchor WLC that becomes unreachable are re-associated with another anchor defined for the WLAN. When this happens, assuming web authentication is being used, the client is redirected to the web portal authentication page and required to re-submit their credentials.


**Note**

Multicast traffic is not supported over guest tunnels, even if multicast is enabled on the Cisco Unified Wireless Network.

Web Portal Authentication

The Cisco Centralized Guest Access solution offers a built-in web portal that is used to solicit guest credentials for authentication and offers simple branding capabilities, along with the ability to display disclaimer or acceptable use policy information (see [Figure 10-5](#)).

Figure 10-5 Controller Web Authentication Page



The web portal page is available on all Cisco WLAN controller platforms and is invoked by default when a WLAN is configured for Layer 3 web policy-based authentication.

If a more customized page is required, administrators have the option of importing and locally storing a customized page. Additionally, if an enterprise wants to use an external web server, the controller can be configured to redirect to it in place of using the internal server. See [Guest Access Configuration](#), [page 10-12](#) for web page configuration guidelines.

User Redirection

As is typical for most web-based authentication systems, in order for guest clients to be redirected to the WLC web authentication page, they must launch a web browser session and attempt to open a destination URL. For redirection to work correctly, the following conditions must be met:

- DNS resolution—The guest access topology must ensure that valid DNS servers are assigned via DHCP and those DNS servers are reachable to users prior to authentication. When a client associates to a web policy WLAN for authentication, all traffic is blocked except DHCP and DNS. Therefore, the DNS servers must be reachable from the anchor controller. Depending on the topology, this might require opening up conduits through a firewall to permit DNS or modifying ACLs on an Internet border router.

**Note**

Clients with static DNS configurations might not work depending on whether their configured DNS servers are reachable from the guest network.

- Resolvable Home Page URL—The home page URL of a guest user must be globally resolvable by DNS. If a user home page is, for example, an internal company home page that cannot be resolved outside of their company intranet, that user is not redirected. In this case, the user must open a URL to a public site such as www.yahoo.com or www.google.com.
- HTTP Port 80—If the home page of a user is resolvable, but connects to a web server on a port other than port 80, they are not redirected. Again, the user is required to open a URL that uses port 80 to be redirected to the WLC web authentication page.

**Note**

In addition to port 80, there is an option to configure one additional port number that the controller can monitor for redirection. The setting is available only through the CLI of the controller:
`<controller_name> config> network web-auth-port <port>.`

Guest Credentials Management

Guest credentials can be created and managed centrally using the management system beginning with release 4.0 and later. A network administrator can create a limited privilege account within the management system that permits lobby ambassador access for the purpose of creating guest credentials. With such an account, the only function a lobby ambassador is permitted to do is create and assign guest credentials to controllers that have web-policy configured WLANs.

As with many configuration tasks within the management system, guest credentials are created using templates. Some of the newer guest user template options and capabilities are:

- There are two types of guest templates: one for scheduling immediate guest access with limited or unlimited lifetime, and the other permits administrators to schedule “future” guest access and offers time of day as well as day of week access restrictions.
- The solution offers administrators the ability to e-mail credentials to guest users. Additionally, when the “schedule” guest template is used, the system automatically e-mails credentials for each new day (interval) that access is offered.
- Guest credentials can be applied to the WLC(s) based on a (guest) WLAN SSID and the management system mapping information: campus/building/floor location or based on a WLAN SSID and a specific controller or list of controllers. The latter method is used when deploying guest access using the guest mobility anchor method as discussed in this chapter.

After a lobby ambassador has created a guest template, it is applied to one or more controllers depending on the guest access topology. Only controllers with a “*web*” *policy-configured WLAN* are listed as a candidate controller to which the template can be applied. This is also true when applying guest templates to controllers based on the management system map location criteria.

Guest credentials, once applied, are stored locally on the (anchor) WLC (under Security > Local Net Users) and remain there until expiration of the “Lifetime” variable as defined in the guest template. If a wireless guest is associated and active when their credentials expire, the WLC stops forwarding traffic and returns to the WEBAUTH_REQD policy state for that user. Unless the guest credentials are re-applied (to the controller), the user is no longer able to access the network.

**Note**

The Lifetime variable associated with guest credentials is independent of the WLAN session timeout variable. If a user remains connected beyond the WLAN session timeout interval, they are de-authenticated. The user is then redirected to the web portal and, assuming their credentials have not expired, must log back in to regain access. To avoid annoying redirects for authentication, the guest WLAN session timeout variable should be set appropriately.

Local Controller Lobby Admin Access

In the event that a centralized management system is not deployed or unavailable, a network administrator can establish a local admin account on the anchor controller, which has only lobby admin privileges. A person who logs in to the controller using the lobby admin account has access to guest user management functions. Configuration options available for local guest management are limited in contrast to the capabilities available through the management system, and include:

- User name
- Generate password
- Administrator assigned password
- Confirm the password
- Lifetime—days:hours:minutes:seconds
- SSID
- Only WLANs configured for Layer 3 web policy authentication are displayed
- Description

Any credentials that may have been applied to the controller by the management system are shown when an admin logs into the controller. A local lobby admin account has privileges to modify or delete any guest credentials that were previously created by the management system. Guest credentials that are created locally on the WLC do not automatically appear in the management system unless the controller’s configuration is updated/refreshed in the management system. Locally created guest credentials that are imported into the management system as a result of a WLC configuration refresh appear as a new guest template that can be edited and re-applied to the WLC.

Guest User Authentication

As previously discussed in [Guest Credentials Management, page 10-9](#), when an administrator uses the management system or a local account on a controller to create guest user credentials, those credentials are stored locally on the controller, which in the case of a centralized guest access topology, would be the anchor controller.

When a wireless guest logs in through the web portal, the controller handles the authentication in the following order:

1. The controller checks its local database for username and password and, if present, grants access.

If no user credentials are found, then:

2. The controller checks to see if an external RADIUS server has been configured for the guest WLAN (under WLAN configuration settings). If so, then the controller creates a RADIUS access-request packet with the user name and password and forwards it to the selected RADIUS server for authentication.

If no specific RADIUS servers have been configured for the guest WLAN:

3. The controller checks its global RADIUS server configuration settings. Any external RADIUS servers configured with the option to authenticate “network” users are queried with the guest user credentials. Otherwise, if no RADIUS servers have “network user” checked, and the user has not authenticated as a result of 1 or 2 above, authentication fails.

**Note**

A RADIUS server can still be used to support network user authentication even if the network user check box is cleared under the WLC Security > AAA > RADIUS settings. However, to do so, a server must then be explicitly selected under the Security > AAA Servers settings of a given WLAN.

External Authentication

WLC and the guest account management (lobby ambassador) capabilities can be used only to create and apply guest user credentials for local authentication on the WLC. However, there may be cases where an enterprise already has an existing guest management /authentication solution deployed as part of a wired guest access or NAC solution. If this is the case, the anchor controller/guest WLAN can be configured to forward web portal authentication to an external RADIUS server, as described in [Guest User Authentication](#).

The default protocol used by the controller to authenticate web users is Password Authentication Protocol (PAP). In the event you are authenticating web users to an external AAA server, be sure to verify the protocols supported by that server. The anchor controller can also be configured to use CHAP or MD5-CHAP for web authentication. The web auth protocol type is configured under the Controller configuration settings of the WLC.

External Authentication using Cisco Secure ACS and Microsoft User Databases

If a guest access deployment is planning to use a Microsoft user database in conjunction with Cisco ACS to authenticate guest users, see the following additional Cisco ACS configuration caveats:

http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.0/installation/guide/windows/postin.html.

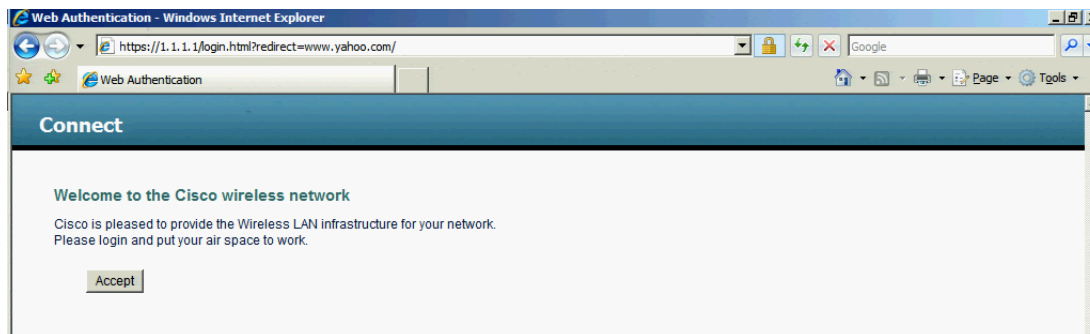
See specifically the following:

http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.0/installation/guide/windows/postin.html#wp1041223

Guest Pass-through

Another variation of wireless guest access is to bypass user authentication altogether and allow open access. However, an enterprise may still need to present an acceptable use policy or disclaimer page to users before granting access. If this is the case, then a guest WLAN can be configured for web policy pass through. In this scenario, a guest user is redirected to a portal page containing disclaimer information.

Pass through mode also has an option for a user to enter an e-mail address before connecting (see [Figure 10-6](#) and [Figure 10-7](#) for sample pages). See [Guest Access Configuration, page 10-12](#) for configuration examples.

Figure 10-6 Pass-through Welcome AUP Page**Figure 10-7** Pass-through Page with E-mail

Guest Users Details E-mail Print Back

Email To:

Subject:

Send Cancel

*Credentials for Guest User **Guest1***

Guest User Name	Guest1
Password	test
Profile	Guest
Start Time	8: 17: 07/19/2007
End Time	9: 0: 07/19/2007

Guest Access Configuration

This section describes how to enable a wireless guest access service within the Cisco Unified Wireless Network solution. The configuration tasks require the use of a web browser. A web session is established with the controller by opening an HTTPS session to the controller management IP address: **https://management_IP** or optionally to a controller service port IP address.

The following procedures assume there is already a deployed infrastructure of controllers and LAPs with the possible exception of the anchor WLC(s). For more information, see: [Anchor Controller Deployment Guidelines](#), page 10-5.



Note

Cisco recommends that the configuration steps outlined in this section be followed in the order in which they are presented.

The following references are used throughout the configuration sections:

- Foreign WLC—Refers to the one or more WLCs deployed throughout an enterprise campus or at branch location that are used for managing and controlling a group of APs. Foreign controllers map a guest WLAN into a guest mobility EoIP tunnel.
- Anchor WLC—Refers to one or more WLCs deployed in the enterprise DMZ that are used to perform guest mobility EoIP tunnel termination, web redirection, and user authentication.



Note

Only the relevant portion of a given configuration screen capture is shown in this section.

The implementation of the Cisco Unified Wireless Network Guest Access solution can be broken into the following configuration categories:

- **Anchor WLC Installation and Interface configuration**—This section briefly discusses installation requirements, steps and caveats associated with implementing one or more anchor WLCs. When implementing guest access for the first time in an existing Cisco Unified Wireless Network deployment, the anchor WLC is usually a new platform that is installed at the Internet edge of an Enterprise network.
- **Mobility Group Configuration**—This section outlines the parameters that must be configured in order for the foreign WLCs to be able to initiate EoIP tunnels to one or more guest anchor WLCs. The mobility group configuration does not itself create the EoIP tunnels, but rather establishes peer relationships between the foreign and anchor WLCs in order to support a guest access WLAN service.
- **Guest WLAN Configuration**—Highlights WLAN specific configuration parameters that are required to map the guest WLAN (originating from a foreign WLC) to the anchor WLC. It is during this portion of the guest access solution configuration that EoIP tunnels are created between the foreign and anchor WLCs. This section also covers the settings required to invoke Layer 3 redirection for web-based authentication.
- **Guest Account Management**—This section outlines how to configure and apply guest user credentials locally on the anchor WLC using controllers the anchor WLC's lobby admin interface.
- **Other Features and Solution Options**—Discusses other features that may be configured including, but not limited to:
 - Web-portal page configuration and management
 - Support for external web redirection
 - Pre-authentication ACLs
 - Anchor WLC DHCP configuration
 - External radius authentication
 - External access control

Anchor WLC Installation and Interface Configuration

As described in [Anchor Controller Positioning, page 10-5](#), Cisco recommends that the anchor WLC be dedicated solely to guest access functions and not be used to control and manage LAPs in the enterprise.

This section does not address all aspects of interface configuration on the anchor WLC. It is assumed the reader is familiar with the WLC initialization and configuration process required upon initial bootup using the serial console interface.

This section offers specific information and caveats as they pertain to configuring interfaces on a WLC being deployed as an anchor in a guest access topology.

As part of the initial configuration (using the serial console interface), you are required to define the following three static interfaces:

- **Controller management**—This interface/IP is used for communications with other controllers in the network. It is also the interface used to terminate EoIP tunnels that originate from the foreign controllers.
- **AP manager interface**—Even though the controller is not used to manage APs, you are still required to configure this interface. Cisco recommends the AP manager interface be configured on the same VLAN and subnet as the management interface.

- Virtual interface—The controller quickstart installation documentation recommends defining the virtual IP with an address, such as 1.1.1.1. This address needs to be the same for all controllers that are members of the same mobility group name. The virtual interface is also used as the source IP address when the controller redirects clients for web authentication.

Guest VLAN Interface Configuration

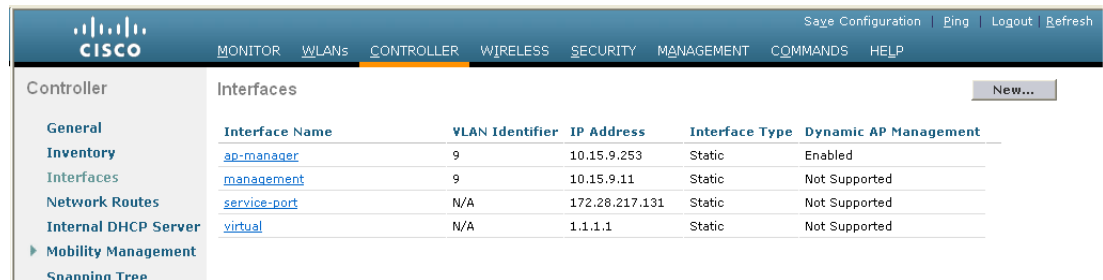
The interfaces previously described are for operations and administrative functions associated with the controller. To implement a guest access service, another interface must be defined. This is the interface through which guest traffic is forwarded for routing to the Internet. As previously described in [Anchor Controller Positioning, page 10-5](#), the guest interface will likely connect to a port on a firewall or be switched to an interface on an Internet border router.

Defining a New Interface

Perform the following to define and configure an interface to support guest traffic:

- Step 1** Click the **Controller** tab.
- Step 2** In the left pane, click **Interfaces** (See [Figure 10-8](#)).

Figure 10-8 *Controller Interfaces*



Interface Name	VLAN Identifier	IP Address	Interface Type	Dynamic AP Management
ap-manager	9	10.15.9.253	Static	Enabled
management	9	10.15.9.11	Static	Not Supported
service-port	N/A	172.28.217.131	Static	Not Supported
virtual	N/A	1.1.1.1	Static	Not Supported

- Step 3** Click **New**.

221855

Defining an Interface Name and VLAN ID

Step 4 Enter an interface name and VLAN ID. (See [Figure 10-9](#).)

Figure 10-9 Interface Name and VLAN ID

The screenshot shows the Cisco Unified Wireless Network Guest Access Services configuration page. The left sidebar contains a navigation menu with options: General, Inventory, Interfaces, Network Routes, Internal DHCP Server, and Mobility Management. The main content area is titled 'Interfaces > New' and contains two input fields: 'Interface Name' with the value 'guest-dmz' and 'VLAN Id' with the value '31'. At the top right of the main area are links for 'Save Configuration', 'Ping', 'Logout', and 'Refresh'. At the bottom right are buttons for '< Back' and 'Apply'.

221856

Defining Interface Properties

Step 5 Define the following properties:

- Interface IP
- Mask
- Gateway (for the firewall or next hop router connected to the anchor controller)
- DHCP Server IP (If using an external DHCP server, use the IP address of that server in the Primary DHCP Server field.)

See [Figure 10-10](#).

Figure 10-10 Defining Interface Properties

The screenshot shows the Cisco Unified Wireless Network Guest Access Services configuration page for defining interface properties. The left sidebar contains a navigation menu with options: General, Inventory, Interfaces, Network Routes, Internal DHCP Server, Mobility Management, Spanning Tree, Ports, Master Controller Mode, Network Time Protocol, QoS, and CDP. The main content area is titled 'Interfaces > Edit' and contains several sections: 'General Information' with fields for 'Interface Name' (guest-dmz) and 'MAC Address' (00:0b:85:40:7e:e0); 'Interface Address' with fields for 'VLAN Identifier' (31), 'IP Address' (10.20.31.11), 'Netmask' (255.255.255.0), and 'Gateway' (10.20.31.1); 'Physical Information' with fields for 'Port Number' (1), 'Backup Port' (0), 'Active Port' (0), and 'Enable Dynamic AP Management' (unchecked); 'Configuration' with a 'Quarantine' checkbox (unchecked); and 'DHCP Information' with fields for 'Primary DHCP Server' (10.20.30.11) and 'Secondary DHCP Server' (empty). At the top right of the main area are links for 'Save Configuration', 'Ping', 'Logout', and 'Refresh'. At the bottom right are buttons for '< Back' and 'Apply'.

221857

**Note**

If DHCP services are to be implemented locally on the anchor controller, populate the primary DHCP server field with the management IP address of the controller. If guest N+1 redundancy is being implemented in the DMZ, repeat the above interface configuration for each additional anchor WLC being deployed.

Mobility Group Configuration

The following default mobility group parameters should already be defined on the foreign WLC(s) as part of a standard centralized WLAN deployment. To support auto-anchor mobility for guest access, the anchor WLC(s) must also be configured with a mobility group domain name.

Defining the Default Mobility Domain Name for the Anchor WLC

Configure a default mobility domain name for the anchor WLC. The anchor's mobility domain name should be different than what is configured for the foreign WLCs. In the examples below, the WLCs (foreign controllers) associated with the enterprise wireless deployment are all members of mobility group 'SRND'. The guest anchor WLC on the other hand, is configured with a different mobility group name: "ANC". This is done to keep the anchor WLC logically separate from the primary mobility domain associated with the enterprise wireless deployment.

- Step 1** Click the Controller tab.
- Step 2** Enter a name in the Default Mobility Domain Name field.
- Step 3** Click **Apply**. (See [Figure 10-11](#).)

Figure 10-11 Defining a Default Mobility Domain Name on the Anchor WLC

The screenshot shows the Cisco Unified Wireless Network configuration interface. The top navigation bar includes tabs for MONITOR, WLANS, CONTROLLER (selected), WIRELESS, SECURITY, MANAGEMENT, COMMANDS, and HELP. The left sidebar shows a tree view with options like General, Inventory, Interfaces, Network Routes, Internal DHCP Server, Mobility Management (selected), Spanning Tree, Ports, Master Controller Mode, Network Time Protocol, QoS, and CDP. The main content area is titled 'Controller' and 'General'. It contains a list of configuration parameters with their current values and dropdown menus for selection. The 'Default Mobility Domain Name' field is highlighted, showing the value 'ANC'. Other fields include RF-Network Name (ANC), User Idle Timeout (300), ARP Timeout (300), Web Radius Authentication (CHAP), and Operating Environment (Commercial (0 to 40 C)). An 'Apply' button is located at the top right of the configuration area.

Parameter	Value	Notes
802.3x Flow Control Mode	Disabled	
LWAPP Transport Mode	Layer 3	(Current Operating Mode is Layer3)
LAG Mode on next reboot	Disabled	(LAG Mode is currently disabled).
Ethernet Multicast Mode	Disabled	
Broadcast Forwarding	Disabled	
Aggressive Load Balancing	Disabled	
Peer to Peer Blocking Mode	Enabled	
Over The Air Provisioning of AP	Disabled	
AP Fallback	Enabled	
Apple Talk Bridging	Disabled	
Fast SSID change	Disabled	
Default Mobility Domain Name	ANC	
RF-Network Name	ANC	
User Idle Timeout (seconds)	300	
ARP Timeout (seconds)	300	
Web Radius Authentication	CHAP	
802.3 Bridging	Disabled	
Operating Environment	Commercial (0 to 40 C)	
Internal Temp Alarm Limits	0 to 65 C	

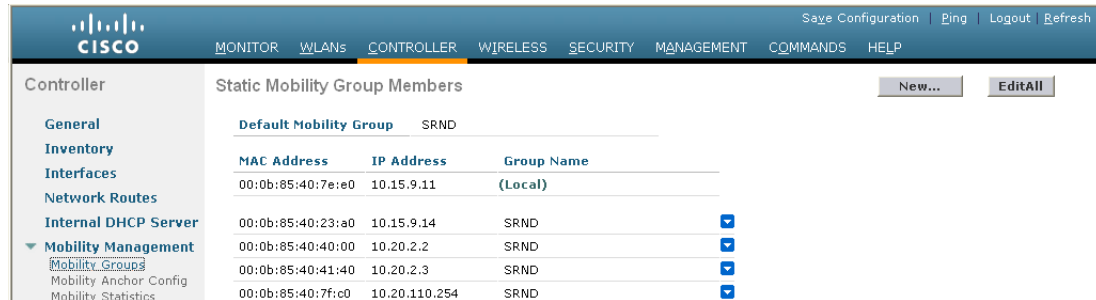
222543

Defining Mobility Group Members of the Anchor WLC

Every foreign WLC within the enterprise deployment that is going to support the guest WLAN must be defined as a mobility group member in the guest anchor WLC(s).

- Step 1** Click the **Controller** tab.
- Step 2** In the left pane, click **Mobility Management** and then **Mobility Groups**. (See [Figure 10-12](#).)

Figure 10-12 Defining Mobility Group Members

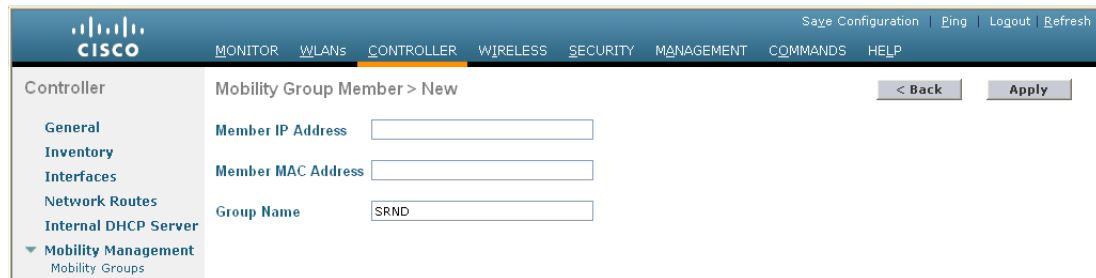


221863

Adding Foreign Controllers as Mobility Group Members

- Step 3** Click **New** to define a MAC and IP address for each foreign controller that will support the guest access WLAN. (See [Figure 10-13](#).)

Figure 10-13 Adding Foreign Controllers to Anchor WLC



221864



Note

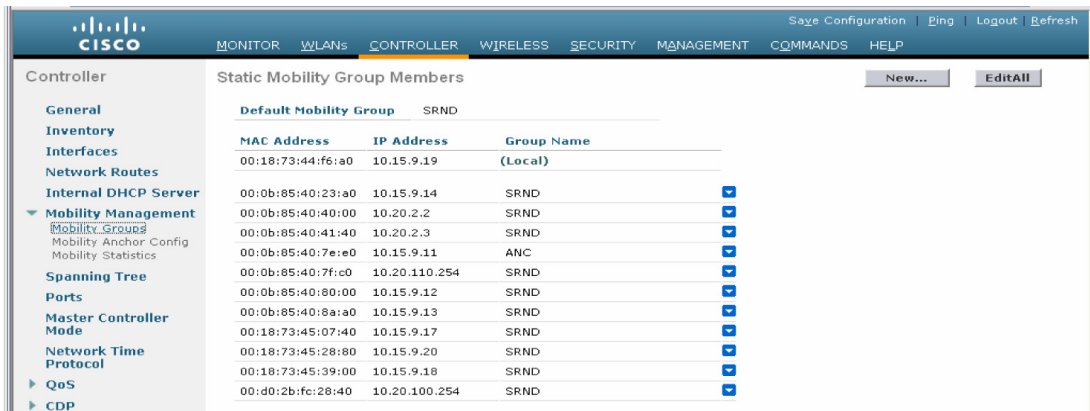
The “Group Name” in [Figure 10-13](#) above is the name configured under the foreign WLC's ‘Default Mobility Domain Name’, which should be different than the name used by the anchor WLC. The member IP and MAC address are those addresses associated with the management interface of the foreign WLCs. Repeat the above steps for each additional foreign WLC that will support the guest WLAN. If more than one anchor is being deployed (guest N+1 redundancy), then repeat the steps in [Defining the Default Mobility Domain Name for the Anchor WLC](#) and [Defining Mobility Group Members of the Anchor WLC](#).

Adding the Anchor WLC as a Mobility Group Member of a Foreign WLC

As described in [Auto Anchor Mobility to Support Wireless Guest Access](#), each foreign WLC maps the guest WLAN into an EoIP tunnel that terminates on the anchor WLC. Therefore, the anchor WLC(s) must be defined as a mobility group member in each foreign controller. In the example below, note that the group name entry for the anchor WLC is 'ANC' (see [Defining Mobility Group Members of the Anchor WLC, page 10-17](#)) whereas the other WLCs that comprise the enterprise wireless deployment are members of the mobility group: 'SRND'.

- Step 1** Click **New** to add the anchor WLC's IP, MAC address, and Group Name to the mobility members table.
- Step 2** Repeat these steps for each additional foreign controller. (See [Figure 10-14](#).)

Figure 10-14 Adding Anchor Controller(s) to Foreign WLC



Static Mobility Group Members		
Default Mobility Group		SRND
MAC Address	IP Address	Group Name
00:18:73:44:f6:a0	10.15.9.19	(Local)
00:0b:85:40:23:a0	10.15.9.14	SRND
00:0b:85:40:40:00	10.20.2.2	SRND
00:0b:85:40:41:40	10.20.2.3	SRND
00:0b:85:40:7e:e0	10.15.9.11	ANC
00:0b:85:40:7f:c0	10.20.110.254	SRND
00:0b:85:40:80:00	10.15.9.12	SRND
00:0b:85:40:8a:a0	10.15.9.13	SRND
00:18:73:45:07:40	10.15.9.17	SRND
00:18:73:45:28:80	10.15.9.20	SRND
00:18:73:45:39:00	10.15.9.18	SRND
00:d0:2b:fc:28:40	10.20.100.254	SRND



Note

If guest N+1 anchor redundancy capability is being deployed, two or more anchor WLC entries are added to each foreign WLC's Mobility Group Members list.

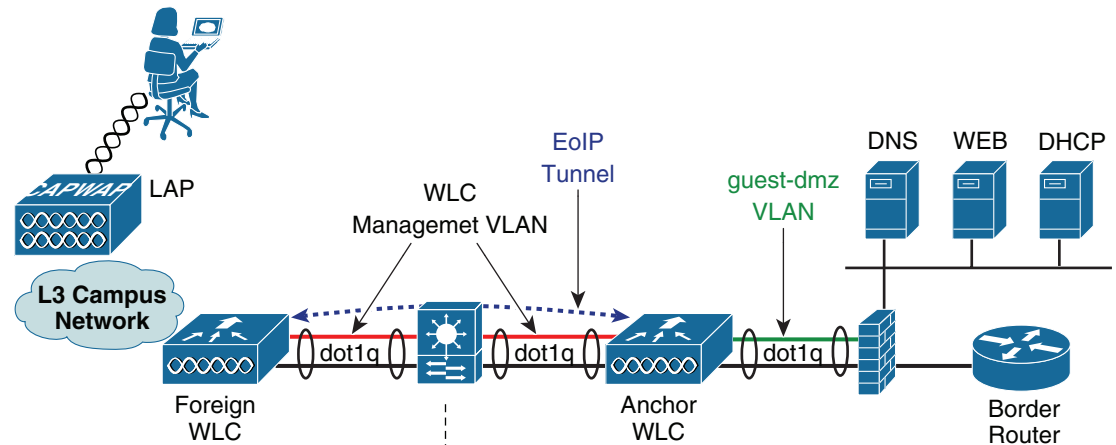
Guest WLAN Configuration

The following section describes how to configure a single guest WLAN. The guest WLAN is configured on every foreign WLC that manages APs where guest access is required. Even though the anchor WLC(s) is not specifically used to manage LAPs associated with a guest WLAN, it must also be configured with the guest WLAN because the anchor WLC is a logical extension of the WLAN where user traffic is ultimately bridged (using CAPWAP between the AP and the foreign controller, and EoIP between the foreign controller and the anchor controller) to an interface/VLAN on the anchor WLC.

**Note**

It is extremely important to note that *all* parameters defined in the WLAN Security, QoS, and Advanced settings tabs, *must be configured identically* in both the anchor and foreign WLC(s). Figure 10-15 shows a high level diagram illustrating the WLAN configuration discussed below.

Figure 10-15 WLAN Configuration



Foreign WLC WLAN Summary

SSID = Guest
 WLAN Status = Enabled
 Radio Policy = 802.11b/g only
 Interface = Management
 Broadcast SSID = Enabled
 Layer 2 Security = None
 Layer 3 Security = None + Web + Auth
 AAA Servers = None
 QOS = Bronze (Background)
 WMM = Disabled
 Advanced = Defaults + DHCP Required

Mobility Config

Default Mobility Group Name = SRND
 Static Mobility Members:
 00:0b:85:40:7e:e0 10.15.9.11 ANC

Anchor WLC WLAN Summary

SSID = Guest
 WLAN Status = Enabled
 Radio Policy = 802.11b/g only
 Interface = guest-dmz
 Broadcast SSID = Enabled
 Layer 2 Security = None
 Layer 3 Security = None + Web + Auth
 AAA Servers = None
 QOS = Bronze (Background)
 WMM = Disabled
 Advanced = Defaults + DHCP Required

Mobility Config

Default Mobility Group Name = ANC
 Static Mobility Members:
 00:18:73:44:f6:a0 10.15.9.19 SRND

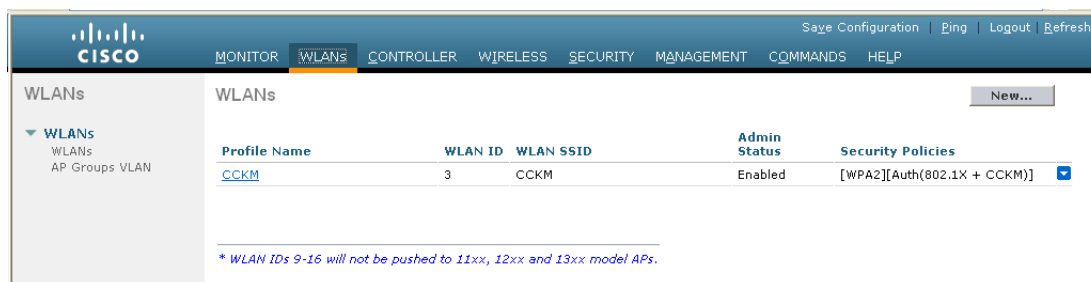
**Note**

The parameters defined in the WLAN Security, QoS, and Advanced settings tabs, *must be configured identically* in both the anchor and foreign controller(s).

Foreign WLC—Guest WLAN Configuration

Step 1 Click the **WLANs** tab and then click **New**. (See [Figure 10-16](#).)

Figure 10-16 Guest WLAN Configuration



221866

Defining a Guest WLAN SSID


Step 2 Define an SSID that is intuitive or easily recognized by potential guest users.

The controller automatically assigns a VLAN ID. Administrators have the option selecting 1 – 16, as long as the ID is not already in use by another SSID/ WLAN.

Step 3 Define a Profile Name.

Step 4 Click **Apply**. (See [Figure 10-17](#).)

Figure 10-17 Defining a Guest WLAN SSID



221867

After creation of the new WLAN, the configuration page appears, as shown in [Figure 10-18](#).

Figure 10-18 WLAN Configuration Page

The screenshot displays the Cisco WLAN Configuration Page. The top navigation bar includes links for MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, and HELP. The left sidebar shows the 'WLANs' menu with sub-items 'WLANs' and 'AP Groups VLAN'. The main content area is titled 'WLANs > Edit' and features tabs for General, Security, QoS, and Advanced. The 'General' tab is active, showing the following configuration details:

- Profile Name: Guest Access WLAN
- WLAN SSID: Guest
- WLAN Status: ☐ Enabled
- Security Policies: [WPA2][Auth(802.1X)] (Modifications done under security tab will appear after applying the changes.)
- Radio Policy: All (dropdown menu)
- Interface: management (dropdown menu)
- Broadcast SSID: ☒ Enabled

Buttons for '< Back' and 'Apply' are located at the top right of the configuration area.



Note

The default interface used by the foreign WLC for the guest WLAN is the management interface. If the EoIP tunnel cannot be established with the anchor, the foreign controller will disassociate any wireless clients that were previously associated with the unreachable anchor and then assign new clients and reassociate clients to the interface configured under the guest WLAN of the foreign itself. Therefore, it is recommended to link the guest WLAN on the foreign to a non-routable network, or alternatively configure the DHCP server of the management interface with an unreachable IP address. If the anchor becomes unreachable, this prevents the guest clients to gain access to the management network.

Defining Guest WLAN Parameters and Policies

Under the General Configuration tab, perform the following steps.

- Step 1** Enable the WLAN by clicking the box next to WLAN Status.
- Step 2** Optionally, set the radio policy if you wish to restrict which bands support the guest access.
 - a. Broadcast SSID is enabled by default; leave enabled.
 - b. By default, the WLAN is assigned to the “management” interface of the WLC. Do not change this.


Step 3 Click the **Security** tab. (See [Figure 10-19](#).)

Figure 10-19 Defining Guest WLAN General Policies



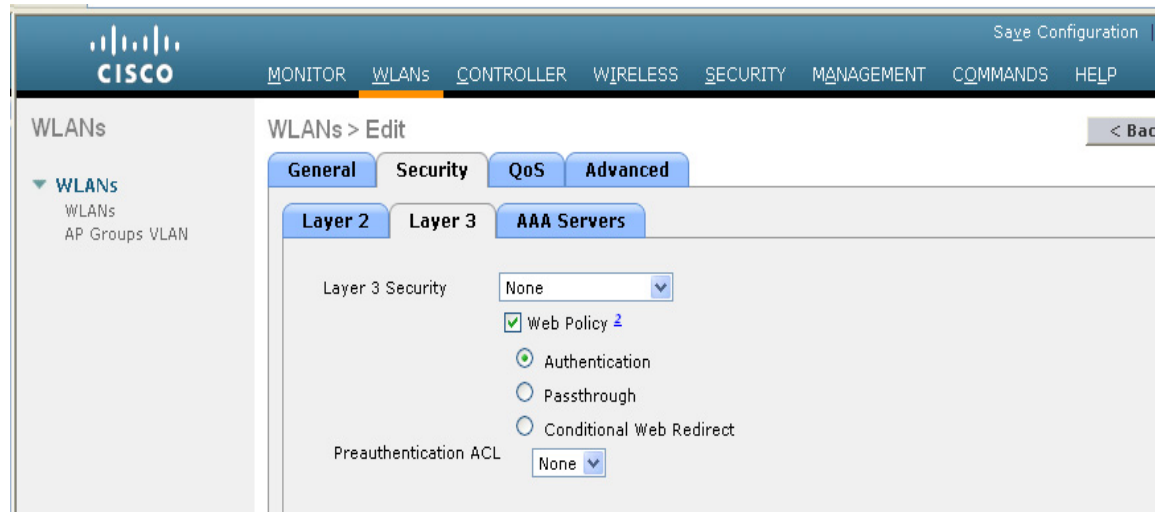
Step 4 Set the Layer 2 Security to **none** from its default setting (802.1x WPA/WPA2). (See [Figure 10-20](#).)

Figure 10-20 WLAN Layer 2 Security Configuration



Step 5 Click the **Layer 3** tab. (See [Figure 10-21](#).)

Figure 10-21 Guest WLAN Layer 3 Security Configuration



Step 6 Click the **Web Policy** checkbox (a list of additional options will be presented).

A dialog warning box appears, indicating that the WLC will pass DNS traffic to and from clients prior to authentication.

Step 7 Select **Authentication** or **Pass-through** for the web policy. (See [Guest User Authentication](#), page 10-10.)

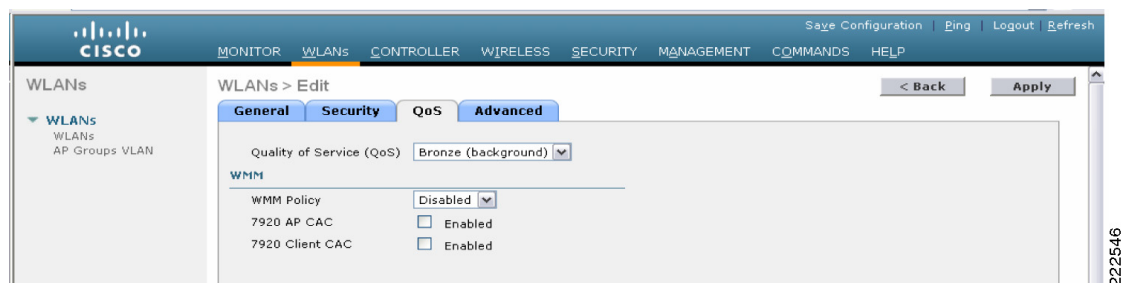


Note

A pre-authentication ACL can be used to apply an ACL that allows un-authenticated clients to connect to specific hosts or URL destinations before authentication. The ACL is configured under Security > Access Control Lists. If a pre-authentication ACL is used in conjunction with the web auth policy, it must include a rule to permit DNS requests; otherwise, the client will be unable to resolve and connect to a destination host/URL that would otherwise be allowed by the ACL.

Step 8 Select the **QoS** tab, as shown in [Figure 10-22](#).

Figure 10-22 Guest WLAN QoS Configuration



Step 9 Optionally, set the upstream QoS profile for the guest WLAN. The default is 'Silver (Best Effort)'. In this example, the guest WLAN has been re-assigned to the lowest QoS class.

Step 10 Click the **Advanced** tab. (See [Figure 10-23](#).)

Figure 10-23 Guest WLAN Advanced Configuration

The screenshot shows the 'WLANs > Edit' configuration page for a specific WLAN. The 'Advanced' tab is active. On the left, there's a sidebar with 'WLANs' and 'AP Groups VLAN'. The main area contains several configuration sections: 'General' (with checkboxes for Allow AAA Override, H-REAP Local Switching, Aironet IE, Diagnostic Channel, and IPv6 Enable), 'Session Timeout (secs)' (a text field with '0'), 'Override Interface ACL' (a dropdown menu set to 'None'), 'Client Exclusion' (checkbox checked, 'Enabled', with a 'Timeout Value (secs)' of '60'), 'DHCP' (checkbox checked, 'Override', 'DHCP Addr. Assignment' set to 'Required'), and 'Management Frame Protection (MFP)' (checkbox checked, 'Infrastructure MFP Protection' set to '(Global MFP Disabled)', 'MFP Client Protection' set to 'Optional'). Buttons for '< Back' and 'Apply' are at the top right.

Step 11 Set Session Timeout (this is optional).



Note Any session timeout greater than 0 (default) forces de-authentication after expiration, and requires the user to re-authenticate through the web portal.

Step 12 Set DHCP Addr. Assignment to “Required”.



Note Setting DHCP Addr. Assignment to “Required” is recommended to prevent guest users from attempting to use the guest network using a static IP configurations.

Step 13 Click **Apply** when finished.

Establishing the Guest WLAN Mobility Anchor(s)

Step 1 From the WLAN menu on the foreign WLC find the newly created guest WLAN.

Step 2 Highlight and click **Mobility Anchors** from the right-hand pull-down selection list. (See [Figure 10-24](#).)

Figure 10-24 WLAN Mobility Anchor

The screenshot shows the 'WLANs' configuration page. A table lists the configured WLANs:

Profile Name	WLAN ID	WLAN SSID	Admin Status	Security Policies
Guest Access WLAN	1	Guest	Enabled	Web-Auth
CCKM	3	CCKM	Enabled	[WPA + WPA2][Auth]

A pull-down menu is open next to the 'CCKM' entry, showing options: 'Remove', 'Mobility Anchors', and a third option (partially visible as 'Remove'). A note at the bottom states: '* WLAN IDs 9-16 will not be pushed to 11xx, 12xx and 13xx model APs.'

Step 3 In the Switch IP Address (Anchor) pull-down selection list, select the IP address corresponding to the management interface of the anchor WLC deployed in the network DMZ. This is the same IP address configured in [Adding the Anchor WLC as a Mobility Group Member of a Foreign WLC](#), page 10-18.

Step 4 Click **Mobility Anchor Create**. (See [Figure 10-26](#).)

Figure 10-25 *Selecting Management Interface from Switch IP Address (Anchor)*

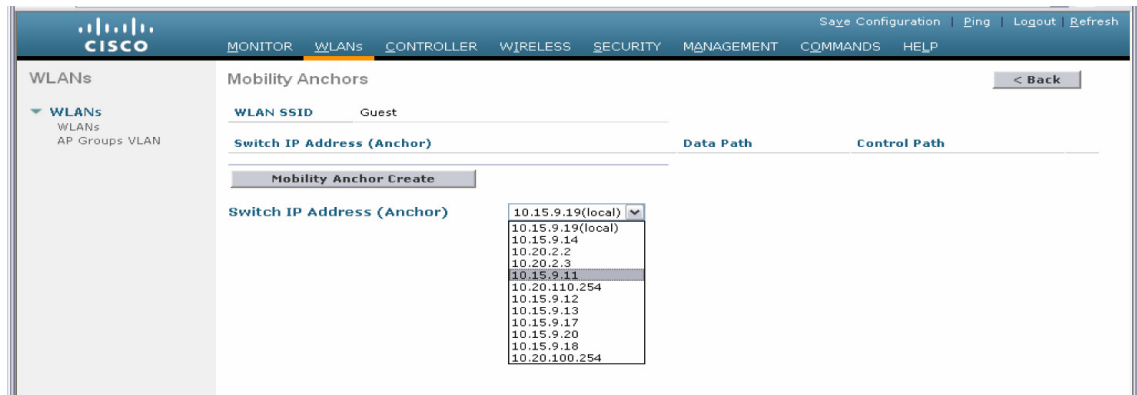
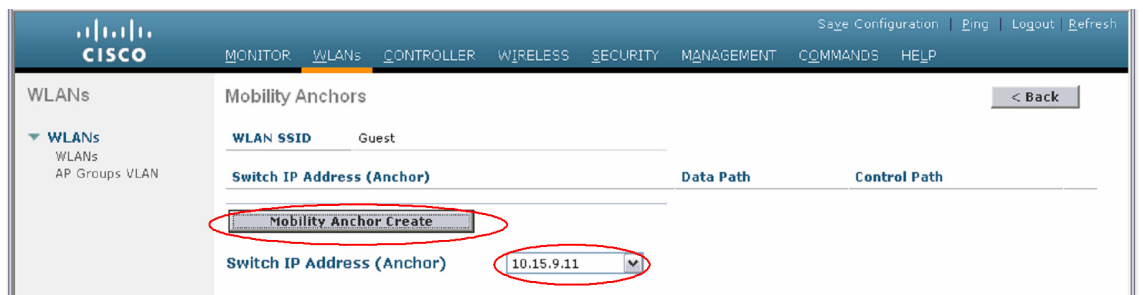


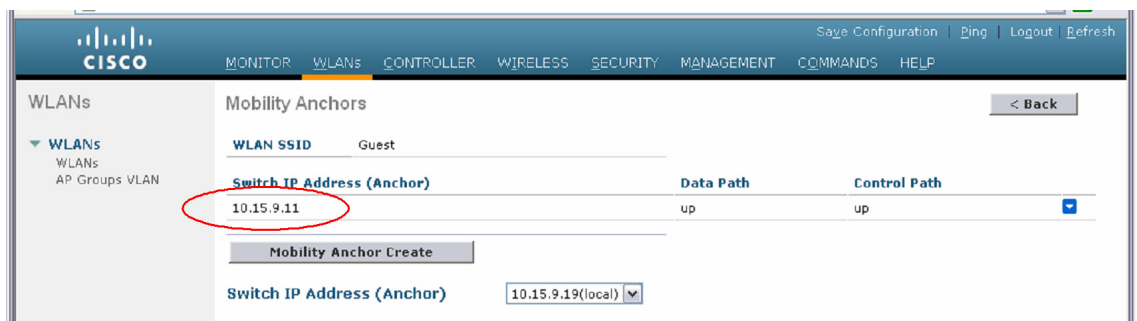
Figure 10-26 *Selecting WLAN Mobility Anchor*



Verifying the Guest WLAN Mobility Anchor

Once configured, the screen shown in [Figure 10-27](#) shows the mobility anchor (selected from above), assigned to the Guest WLAN.

Figure 10-27 *Verifying the Guest WLAN Mobility Anchor*



For ease of verification, the page displays whether or not the mobility tunnel data path and CAPWAP control path have been established with the anchor. If either or both show “down”, see [Troubleshooting Guest Access, page 10-56](#) for troubleshooting tips. The pull-down selection list to the right offers the option to send a ping to the destination anchor WLC.

- Step 5** When finished, click **Back**.
- Step 6** Repeat the steps above for each additional anchor WLC being deployed (guest N+1 redundancy).

This completes the guest WLAN configuration. Repeat all steps from [Foreign WLC—Guest WLAN Configuration](#) through [Verifying the Guest WLAN Mobility Anchor](#) for each additional foreign WLC that will support the guest WLAN.

Guest WLAN Configuration on the Anchor WLC

Guest WLAN configuration on the anchor controller(s) is identical to that of the foreign controller except for minor differences in the WLAN interface and mobility anchor configuration, which are detailed below.



Note

The SSID defined for the guest WLAN must be exactly the same as what is defined on the foreign WLCs.

Anchor WLC—Guest WLAN Interface

As indicated above, the parameters configured for the guest WLAN on the anchor WLC are the same except the interface to which the WLAN is mapped. In this case, the guest WLAN is assigned to an interface/VLAN on the anchor WLC, which connects to an interface on a firewall or Internet border router.

- Step 1** Click the **WLANs** tab.
- Step 2** Create, configure, and enable the guest WLAN the same way it was configured on the foreign WLC(s) except for the following:
- In the WLANs general configuration, under **Interface**, choose the interface name created in [Guest VLAN Interface Configuration](#). (See [Figure 10-28](#).)
- Step 3** Click **Apply**.

Figure 10-28 Anchor WLC Guest WLAN Interface Configuration

The screenshot shows the Cisco WLC configuration interface. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh. The main menu has tabs for MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, and HELP. The left sidebar shows a tree view with WLANs expanded, showing AP Groups and VLAN. The main content area is titled 'WLANs > Edit' and has tabs for General, Security, QoS, and Advanced. The General tab is selected, showing the following configuration:

- Profile Name: Guest Access WLAN
- WLAN SSID: Guest
- WLAN Status: ☒ Enabled
- Security Policies: [WPA2][Auth(802.1X)] (Modifications done under security tab will appear after applying the changes.)
- Radio Policy: 802.11b/g only
- Interface: guest-dmz (highlighted with a red circle)
- Broadcast SSID: ☒ Enabled

Buttons for '< Back' and 'Apply' are visible at the top right of the configuration area.

Anchor WLC—Defining the Guest WLAN Mobility Anchor

The second parameter that differs in configuration from the foreign WLC is the WLAN mobility anchor configuration. The guest WLAN mobility anchor is the anchor WLC itself.

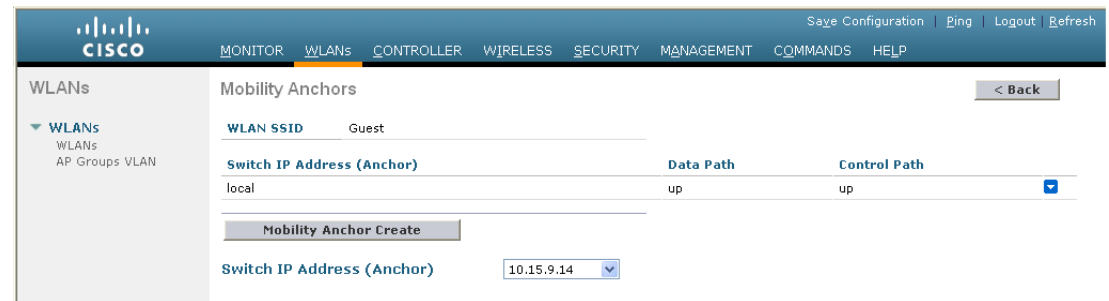
- Step 1** Click the **WLANs** tab.
- Step 2** Find the Guest WLAN and click **Mobility Anchors**.
- Step 3** From the pull-down selection list, choose the IP address representing the anchor controller. The IP address has (Local) next to it.
- Step 4** Click **Mobility Anchor Create**. (See [Figure 10-29](#).)

Figure 10-29 Defining the Guest WLAN Mobility Anchor



Note that the guest WLAN mobility anchor is *local*. (See [Figure 10-30](#).)

Figure 10-30 Verifying Guest Mobility Anchor



Because the mobility anchor for the guest WLAN is the anchor WLC itself, the Data and Control Path status will always show “up”. If not, check to ensure that you have selected the local WLC as the anchor from the 'Switch IP Address (Anchor)' drop down menu.

- Step 5** If guest N+1 redundancy is being implemented, repeat the WLAN configuration for each additional anchor WLC being deployed. Otherwise, this completes the configuration steps required to create the guest WLAN on the anchor WLC.

Guest Account Management

- If guest credentials are going to be managed locally on the anchor controller, there are two methods by which they can be created and applied:
- Through a lobby ambassador admin or super user/root admin account
- Directly on the controller via a local lobby admin account or other management account with read/write access

Guest Management Using the Management System

The following configuration examples assume the management system version 4.1.83 or later has been installed and configured, and a lobby ambassador account has been created.

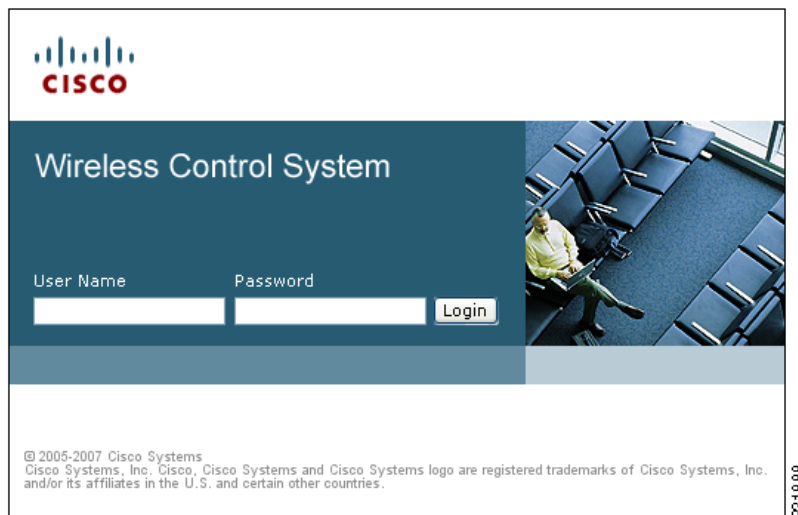


Note

Ensure that the individual WLC configurations are synchronized with the management system before creating guest templates.


Log in to the management system using the Lobby Ambassador credentials assigned by the system administrator. (See [Figure 10-31](#).)

Figure 10-31 Lobby Ambassador



After logging in, the screen shown in [Figure 10-32](#) appears.

Figure 10-32 Cisco Prime Infrastructure Lobby Admin Interface




Note

Cisco Prime Infrastructure was formally known as WCS and NCS.

There are two types of guest templates:

- The **Add Guest User** template allows administrators to create and immediately apply guest credentials to one or more anchor WLCs.
- The **Schedule Guest User** template allows administrators to create guest credentials that are applied to one or more anchor WLCs at some future month, day, and time. (See [Figure 10-33](#).)

Figure 10-33 Guest User Template Option



Using the Add Guest User Template

- Step 1** From the pull-down selection list, select **Add Guest User** and click **Go**.
- Step 2** The template shown in [Figure 10-34](#) appears.

Figure 10-34 Add Guest User Template

The screenshot shows the Cisco Wireless Control System interface for adding a new guest user. The page is titled "Guest Users > New User" and includes a sidebar with "Guest Users" and a top navigation bar with "Help". The main content area is divided into two sections: "Guest Information" and "Account Configuration".

Guest Information:

- User Name:
- Generate Password: ☐
- Password:
- Confirm Password:

Account Configuration:

- Profile:
- Life Time: ☒ Limited ☐ Unlimited
- End Time: Hour Min. Day
- Apply To:
- Campus:
- Building:
- Floor:
- Description:
- Disclaimer:
- ☐ Make this Disclaimer default

At the bottom, there are "Save" and "Cancel" buttons. The page number "221891" is visible in the bottom right corner.

Figure 10-35 shows an example of guest user account creation.

Figure 10-35 Guest User Account Creation

Step 3 Under **Guest Information**, enter a User Name and Password.

Passwords are case sensitive. User names are restricted to 24 characters or less. Administrators also have an option to allow the system to automatically generate a password by clicking on the **Generate Password** check box.

Step 4 Under **Account Configuration**, select the following:

- Profile—The pull-down selection list displays a list of WLANs (SSIDs) configured with a L3 Web Policy.
- Life Time—Select “limited” or “unlimited”
- End Time—If the guest account is “limited”, select the month, day, and time the credentials are to expire.
- Apply To—From the pull-down selection list, select **Controller List** and click the check box next to the controller(s) representing anchor WLCs. Note that there will be other controllers listed; however, these represent the foreign WLCs. There is no need to apply user credentials on the foreign WLCs because the authentication enforcement point is the anchor WLC.



Note

As seen in Figure 10-35, there are various options for where the credentials can be applied, including being able to control the physical/geographic location where a user can access the guest WLAN. These include outdoor areas, indoor areas, building, floor, and so on. This location-based access method can only be used if: 1) the WLAN deployment has been integrated into the management system mapping database, and 2) the guest WLAN (a WLAN with web policy) does not use mobility anchors.

- **Description**—Enter a description. The description is displayed on the WLC to which the credentials are applied under Security > Local Net Users. It is also included in the e-mail that can be sent to a guest informing them of what credentials to use to access the network.
- **Disclaimer**—Used in the e-mail that can be sent to a guest user informing them of what credentials to use to access the network

Step 5 Click **Save** when finished. The summary screen shown in [Figure 10-36](#) appears, acknowledging that credentials have been applied to the anchor controller(s). The admin is also presented with an option to print or e-mail the credentials to the guest user.

Figure 10-36 Successful Guest Account Creation

The screenshot shows the Cisco Wireless Control System (WCS) interface. The top navigation bar includes the Cisco logo, the title 'Wireless Control System', and user information: 'Username: lobbyadmin | Logout | Refresh | Print View'. A 'Help' dropdown menu is visible. The left sidebar has a 'Guest Users' section. The main content area displays the 'Guest User Account application result to the Selected controllers'.

IP Address	Controller Name	Operation Status	Reason
10.15.9.11	Controller1	Success	-
10.15.9.13	Controller3	Success	-

Below the table, there is a section titled 'Guest User Credentials' with the following details:

Guest User Name	Guest1
Password	test
Profile	Guest
Start Time	8: 17: 07/19/2007
End Time	9: 0: 07/19/2007
Disclaimer	Guests understand and acknowledge that we exercise no control over the nature, content or reliability of the information and/or data passing through our network.

At the bottom of the credentials section, there is a link: [Print/Email Guest User Credentials](#).

221893

Step 6 Click **Print/Email Guest User Credentials**. The screen shown in [Figure 10-37](#) appears.

Figure 10-37 *Print/Email Guest User Details*

Guest Users Details	
Email To	<input type="text"/>
Subject	<input type="text"/>
<input type="button" value="Send"/> <input type="button" value="Cancel"/>	
Credentials for Guest User Guest1	
Guest User Name	Guest1
Password	test
Profile	Guest
Start Time	8: 17: 07/19/2007
End Time	9: 0: 07/19/2007
Disclaimer	Guests understand and acknowledge that we exercise no control over the nature, content or reliability of the information and/or data passing through our network.



Note

For details on setting up an SMTP mail server to support e-mailing guest account information to users, see the WCS Configuration guide at:

<http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/WCS60cg.html>.

After printing and or e-mailing the account details, the screen shown in [Figure 10-38](#) appears. By clicking the **User Name**, an admin can go back and edit the guest account or remove it by checking the box next to the User Name and selecting **Delete Guest User** from the pull-down selection list.

Figure 10-38 *Cisco Prime Infrastructure Guest Users Summary*

User Name	Profile	Description	Applied To	Status
<input type="checkbox"/> Guest1	Guest	Wireless Network Guest Access	Controller List	Active



Note


If a user template is deleted from Cisco Prime Infrastructure while a user is active, they are de-authenticated.

Using the Schedule Guest User Template

For details about configuring guest accounts, see Cisco Wireless Control System Configuration Guide at: <http://www.cisco.com/en/US/docs/wireless/wcs/4.1/configuration/guide/wcsadmin.html>

Figure 10-39 shows the guest user template option.

Figure 10-39 Guest User Template Option



- Step 1** From the pull-down selection list, select **Schedule Guest User** and click **Go**.
The template shown in Figure 10-40 appears.

Figure 10-40 Schedule Guest User Template

Wireless Control System Username: lobbyadmin | Logout | Refresh | Print View

Help ▾

Guest Users

Guest Users > Scheduling

Guest Information

User Name

☐ Generate new on every schedule

Account Configuration

Profile

Life Time ☒ Limited ☐ Unlimited

Start Time (Hours) (Minutes)

End Time (Hours) (Minutes)

Days of the week ☐ Sun ☐ Mon ☐ Tues ☐ Wed ☐ Thur ☐ Fri ☐ Sat

Apply to

Campus

Building

Floor

Email credentials to

Description

Disclaimer

☐ Make this Disclaimer default

Save **Cancel**

Figure 10-41 shows an example of a schedule guest user account creation.

Figure 10-41 Schedule Guest User Account Creation

The screenshot displays the 'Wireless Control System' interface for 'Guest Users > Scheduling'. The 'Guest Information' section includes a 'User Name' field with 'test2' and an unchecked checkbox for 'Generate new on every schedule'. The 'Account Configuration' section features a 'Profile' dropdown set to 'Guest', 'Life Time' radio buttons for 'Limited' (selected) and 'Unlimited', 'Start Time' (8:00 on 07/19/07), 'End Time' (17:00 on 07/27/07), and 'Days of the week' checkboxes for Sun, Mon, Tues, Wed, Thur, Fri, and Sat. Below this is a 'Controller List' table with columns for 'IP Address' and 'Name', showing three entries: 10.15.9.11 (Controller1), 10.15.9.13 (Controller3), and 10.15.9.19 (Controller9). Other fields include 'Email credentials to' (john.doe@crisco.com), 'Description' (Wireless Network Guest Ac), and a 'Disclaimer' text area. 'Save' and 'Cancel' buttons are at the bottom.

Step 2 Under Guest Information, enter a User Name. User names can be up to 24 characters long. When using the schedule-based template, administrators have the option to allow the system to automatically generate the user name for each new day that access is being offered. Also, when using this template, the system automatically generates the user password. There is no option to manually assign a password.

Step 3 Under Account Configuration, select the following:

- Profile—The pull-down selection list displays a list of WLANs (SSIDs) configured with an L3 Web Policy.
- Life Time—Select “limited” or “unlimited”.
- Start Time—Select the time, month, and day when the account is to become active.



Note The start time cannot begin within the current day that the account is being created. The start day must be one or more days beyond the day the account is being created.

- End Time—If the account is limited, select the stop time, month, and day.



Note The stop day can be a period no longer than 30 days from the start day.

- Days of Week—Depending on the lifetime of the account, administrators have the ability to control for which days of the week access is available. Click the check boxes next to those days of the week access is permitted.

**Note**

If “Days of the Week” is selected, the start and stop times represent the period within each day that access is available. Upon expiry within a given day, Cisco Prime Infrastructure removes the credentials from the applicable controllers. For each new day/interval that access is permitted, Cisco Prime Infrastructure automatically generates a new password (and optionally a username), e-mails it to the guest user, and re-applies the new credentials to the applicable WLCs. If “Days of the Week” is not defined, access begins based on the start day and time and is continuously active until the end day and time.

- **Apply To**—From the pull-down selection list, select **Controller List** and click the check box next to the controller(s) representing anchor WLCs. Note that there will be other controllers listed; however, these represent the foreign WLCs. There is no need to apply user credentials on the foreign WLCs because the authentication enforcement point is the anchor WLC.

**Note**

As seen in [Figure 10-41](#), there are various options for where the credentials can be applied, including being able to control the physical/geographic location where a user can access a guest WLAN. These include outdoor areas, indoor areas, building, floor, and so on. This location-based access method can only be used if: 1) the WLAN deployment has been integrated into the Cisco Prime Infrastructure mapping database, and 2) the guest WLAN (a WLAN with web policy) does not use mobility anchors.

- **E-mail Credentials to**—Enter the e-mail address for whom an account is being established. This is a mandatory field.

**Note**

An SMTP mail server must be configured in Cisco Prime Infrastructure so that it can use to send guest account information. For details, see:
http://www.cisco.com/en/US/docs/wireless/wcs/6.0/configuration/guide/6_0admin.html.

- **Description**—Provide a description. The description is displayed on the WLC to which the credentials are applied under Security > Local Net Users. The description is also included in an e-mail that can be sent to the guest, informing them of what credentials to use to access the network.
- **Disclaimer**—Used in the e-mail that is sent to a guest user, informing them of what credentials to use to access the network.

- Step 4** Click **Save** when finished. The screen shown in [Figure 10-42](#) appears, acknowledging that the scheduled account has been created. The admin is also presented with an option to print or e-mail the credentials to the guest user.

Figure 10-42 Successful Scheduled Account Creation

Wireless Control System Username: lobbyadmin | Logout | Refresh | Print View

Guest Users

Guest User Account Scheduled on the selected controllers

Guest User Credentials

Guest User Name	test2
Password	Frla4urF
Profile	Guest
Start Time	8: 0: 07/20/2007
End Time	17: 0: 08/03/2007
Disclaimer	Guests understand and acknowledge that we exercise no control over the nature, content or reliability of the information and/or data passing through our network.

[Print/Email Guest User Credentials](#)

- Step 5** Optionally, click **Print/Email Guest User Credentials**. The screen shown in [Figure 10-43](#) appears.

Figure 10-43 Print/E-mail Guest User Details

Guest Users Details E-mail Print Back

Email To

Subject

Send Cancel

Credentials for Guest User test2

Guest User Name	test2
Password	Frla4urF
Profile	Guest
Start Time	8: 0: 07/20/2007
End Time	17: 0: 08/03/2007
Disclaimer	Guests understand and acknowledge that we exercise no control over the nature, content or reliability of the information and/or data passing through our network.

After printing and/or e-mailing the account details, the summary screen shown in [Figure 10-44](#) appears. By clicking the **User Name**, an admin can go back and edit the guest account or remove it by checking the box next to the User Name and selecting **Delete Guest User** from the pull-down selection list.

Figure 10-44 Cisco Prime Infrastructure Guest Users Summary

Wireless Control System Username: lobbyadmin | Logout | Refresh | Print View

Guest Users Add Guest User GO

<input type="checkbox"/> User Name	<input type="checkbox"/> Profile	Description	Applied To	Status
<input type="checkbox"/> test2	<input type="checkbox"/> Guest	Wireless Network Guest Access	Controller List	Scheduled



Note

If a user template is deleted from Cisco Prime Infrastructure while a user is active, they are de-authenticated.

This completes the steps required to create a guest account using the lobby ambassador interface in Cisco Prime Infrastructure.


Managing Guest Credentials Directly on the Anchor Controller

The following procedure assumes that a network administrator has established a local management account with lobby admin privileges on one or more anchor controllers.

- Step 1** Login to the anchor controller using the lobby admin credentials assigned by the system administrator. Remember that conduits might need to be opened through a firewall to permit HTTP/HTTPS for web administration of the controller. See [Anchor Controller Positioning, page 10-5](#).

After login, the screen shown in [Figure 10-53](#) appears.

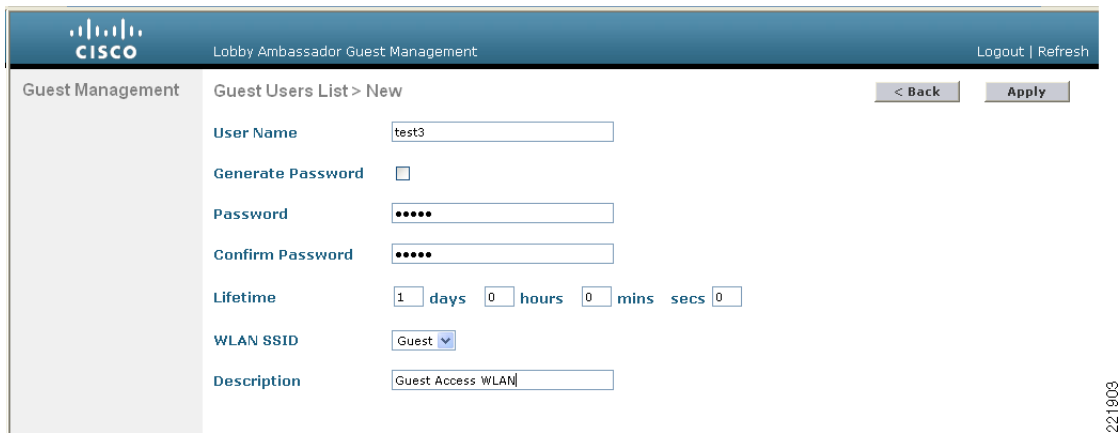
Figure 10-45 *Anchor Controller Login*



221902

- Step 2** Click New.
The screen shown in [Figure 10-46](#) appears.

Figure 10-46 *Creating Local WLC Guest Credentials*



221903

- Step 3** To create user credentials, perform the following steps:
- a. Enter a username and password (manual or auto).

- b. Select the WLAN/SSID to which the guest account applies (only WLANs configured with an L3 web policy are displayed).
- c. Enter a lifetime for the credentials.
- d. Enter a description for the user.

Step 4 Click **Apply**.

The screen shown in [Figure 10-47](#) appears and shows the newly-added guest user.

Figure 10-47 Anchor WLC Guest Users List

User Name	WLAN SSID	Account Remaining Time	Description
test3	Guest	1 d	Guest Access WLAN

From this screen you have the option to do the following:

- Edit the existing user (link at far right; not visible)
- Delete the existing user (link at far right; not visible)
- Add a new user

Configuring the Maximum Number of User Accounts

The default number of guest user accounts that can be defined on the controller is 512. This value can be changed by completing the following steps.

Step 1 Click the **Security** tab. (See [Figure 10-48](#).)

Figure 10-48 Configuring the Maximum Number of User Accounts

Step 2 In the left pane, click **General** under AAA properties.

Step 3 Configure the maximum number of user database entries (between 512 and 2048).

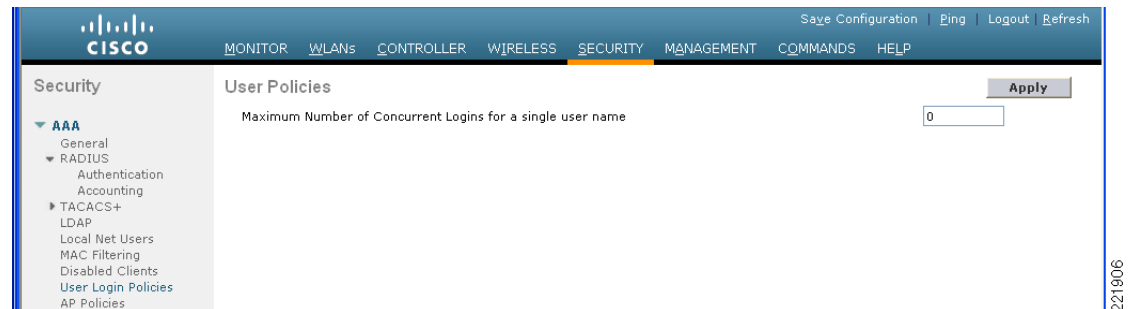
Step 4 Click **Apply**.

Maximum Concurrent User Logins

The maximum number of concurrent logins for a local user account on the WLC can be configured. Values include 0 for unlimited concurrent logins or can be limited from 1 to 8. The maximum user logins is configured by completing the following steps:

Step 1 Click the **Security** tab. (See [Figure 10-49](#).)

Figure 10-49 User Login Policies



Step 2 In the left pane, click **User Login Policies** under AAA.

Step 3 Configure the maximum number of concurrent user logins (between 0–8).

Step 4 Click **Apply**.

Guest User Management Caveats

Note the following caveats:

- Guest accounts can be added using either method above or both methods together.
- When using Cisco Prime Infrastructure, the lobby admin may not have visibility of user accounts that might have been created locally on the anchor controller if the controller configuration has not been recently synchronized with Cisco Prime Infrastructure. If this is the case and a Cisco Prime Infrastructure lobby admin attempts to add an account with a user name that is already configured on the WLC, the Cisco Prime Infrastructure configuration overrides the local configuration.
- When adding user accounts locally on the controller, the local admin will have visibility of all accounts that have been created, including those that were created via Cisco Prime Infrastructure.
- If a guest user is currently authenticated to a WLAN and their credentials are deleted from Cisco Prime Infrastructure or locally on the controller, the user traffic stops flowing, and the user is de-authenticated.

Other Features and Solution Options

Web Portal Page Configuration and Management

The internal web server and associated functionality is hosted locally on the anchor controller. When a WLAN is configured to use the web policy, either for authentication or pass-through, the internal web server is invoked by default. No further configuration is required. The internal portal includes a few optional configuration parameters.

Internal Web Page Management

Step 1 Click the **Security** tab.

Step 2 In the left pane, click **Web Auth** and then **Web Login Page**.

The configuration screen shown [Figure 10-50](#) is displayed. You can change the heading and message information that appears on the portal page. You can also choose a post-authentication redirect URL.

Figure 10-50 Web Login Page Configuration Screen

The screenshot displays the Cisco Web Login Page Configuration interface. The top navigation bar includes tabs for MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY (selected), MANAGEMENT, COMMANDS, and HELP. The left sidebar shows the Security configuration tree, with 'Web Auth' and 'Web Login Page' highlighted. The main configuration area for the 'Web Login Page' includes:

- Web Authentication Type:** A dropdown menu set to 'Internal (Default)'.
- Description:** A text block explaining that this page allows customization of the login page content and appearance.
- Cisco Logo:** Radio buttons for 'Show' (selected) and 'Hide'.
- Redirect URL after login:** An empty text input field.
- Headline:** A text input field containing 'Welcome to the Cisco wireless network'.
- Message:** A large text area containing a sample message: 'Cisco is pleased to provide the Wireless LAN infrastructure for your network. Please login and put your air space to work.'

Buttons for 'Preview...' and 'Apply' are located at the top right of the configuration area.

221883

Step 3 Click **Apply**.

Step 4 Optionally, click **Preview** to view what the user sees when redirected.

Importing a Web Page

You can download a customized web page and store it locally on the anchor controller. To import a customized web page, perform the following steps.

Step 1 Click the **Commands** tab. (See Figure 10-51.)

Figure 10-51 Importing a Web Page

The screenshot shows the Cisco WLC web interface. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh. The main menu has tabs for MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS (selected), and HELP. On the left, the 'Commands' sidebar lists options: Download File, Upload File, Reboot, Reset to Factory Default, and Set Time. The main content area is titled 'Download file to Controller' and features a 'Clear' button and a 'Download' button. A 'File Type' dropdown menu is set to 'Webauth Bundle'. Below this is the 'TFTP Server' configuration section with the following fields: IP Address (10.20.30.200), Maximum retries (10), Timeout (seconds) (6), File Path (/), and File Name (empty).

Step 2 Under File Type, select **Web Auth Bundle**.

Step 3 Define the IP address and file path on the TFTP server where the files reside.

Step 4 Click **Download** to begin.

Be aware of these caveats when downloading a web auth bundle:

- Select **Web Auth Bundle** from the pull-down selection list to ensure that the files are stored in the correct directory on the controller.
- The **Web Auth Bundle** must be a **.tar** file of the HTML and image files associated with the custom web login page. When downloaded, the WLC un-tars the files and places them in the appropriate directory.
- The **Web Auth Bundle** (**.tar** file) cannot be larger than 1 MB.
- The file name for the HTML login page must be **login.html**.

For more information about downloading and using customized web pages, see:

<http://www.cisco.com/en/US/docs/wireless/wcs/4.1/configuration/guide/wcssol.html#wp1065703>.

Selecting an Imported Web Auth Page

To use a customized web auth page that has been downloaded to the controller, perform the following steps:

Step 1 Click the **Security** tab.


Step 2 In the left pane, click **Web Auth** and then **Web Login Page**.

Step 3 From the Web Authentication Type pull-down selection list, select **Customized** (Downloaded).

Step 4 Click **Preview** to view the downloaded page.

Step 5 Click **Apply** when finished. (See [Figure 10-52](#).)

Figure 10-52 *Selecting an Imported Web Auth Page*



221865

Internal Web Certificate Management

The web auth login page uses SSL for safeguarding user credentials. For simplicity, the controller uses a self-signed certificate. Because the certificate is self-signed, guest users can expect to see a pop-up alert similar to the following when they are redirected to the authentication page shown in [Figure 10-53](#).

Figure 10-53 *Web Certificate Security Alert (IE6)*



190842

At this point, you can proceed by either clicking **Yes** or you can select **View Certificate** and manually install it as a trusted site. The web server uses the virtual interface IP address configured in [Anchor WLC Installation and Interface Configuration, page 10-13](#), as its source address. If a hostname is defined along with the IP address, that host name must be resolvable by DNS so that:

- The client is redirected to the web auth page.
- The user does not encounter a web certificate error because of conflicts between hostname and host IP address.

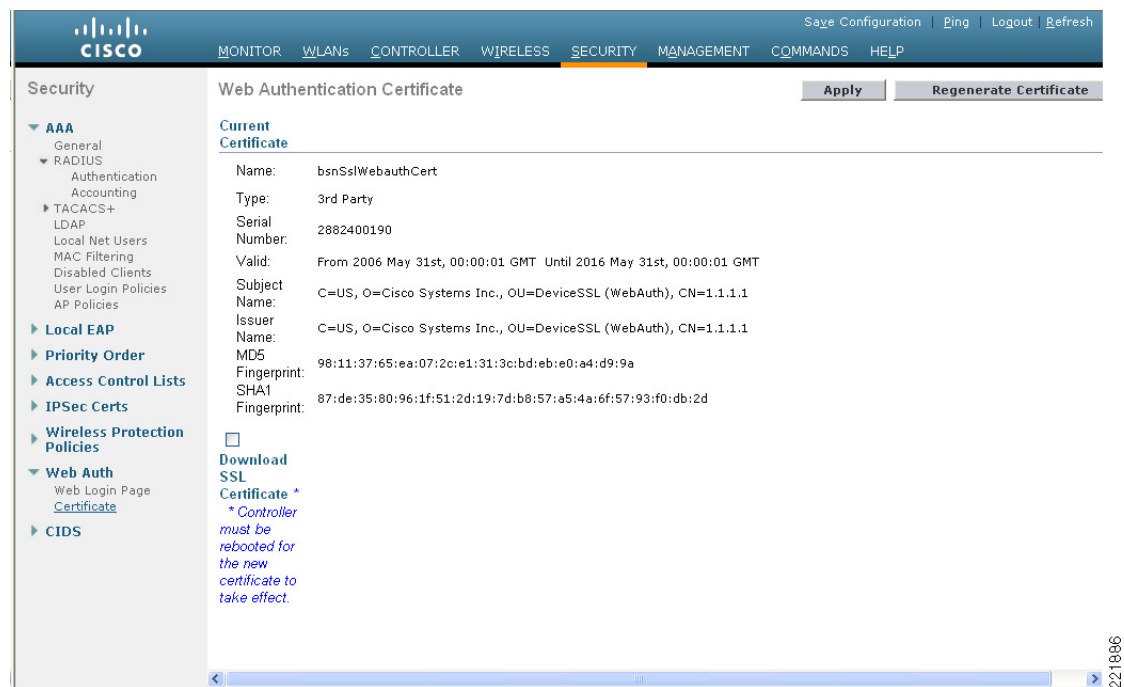
Importing an External Web Certificate

For cases where a legitimate web certificate issued by a trusted root CA is required, one can be downloaded to the controller by performing the following steps:

Step 1 Click the **Security** tab.

In the left pane, click **Web Auth** and then **Certificate**. (See [Figure 10-54](#).)

Figure 10-54 Importing an External Web Certificate



Step 2 Place a check mark in the **Download SSL Certificate** check box.

Step 3 Complete the required fields for downloading the certificate.

Step 4 Click **Apply**.

Step 5 After the certificate has been downloaded, reboot the server.

Support for External Web Redirection

In some cases, an enterprise might already have deployed a web-portal system to support wired guest access or NAC functionality. If this is the case, the anchor controller can be configured to redirect wireless guest users to an external web portal using the following steps:

- Step 1** Click the **Security** tab.
- Step 2** In the left pane, click **Web Auth** and then **Web Login Page**. (See [Figure 10-55](#).)


Figure 10-55 Supporting External Web Redirection

The screenshot shows the Cisco Unified Wireless Network configuration interface. The top navigation bar includes tabs for MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY (selected), MANAGEMENT, COMMANDS, and HELP. The left sidebar shows the Security configuration tree with options like AAA, RADIUS, TACACS+, Local EAP, Priority Order, Access Control Lists, IPSec Certs, Wireless Protection Policies, Web Auth (selected), and CIDS. Under Web Auth, the Web Login Page is selected. The main configuration area for the Web Login Page shows the Web Authentication Type set to 'External (Redirect to external server)'. The URL field is populated with 'https://10.20.30.41'. Below this, there is a section for External Web Servers with a Web Server IP Address field and an 'Add Web Server' button. The interface also includes 'Preview...' and 'Apply' buttons at the top right of the configuration area.


- Step 3** Fill in the Web Server IP and URL fields.
- Step 4** Click **Apply**.

Anchor WLC-Pre-Authentication ACL

A pre-authentication ACL (pre-auth ACL) can be applied to the guest WLAN, which allows unauthenticated clients to connect to specific hosts or URL destinations prior to authenticating. The pre-auth ACL is applied under the guest WLAN Layer 3 Security settings and, if enabled, is performed only on the anchor WLC(s). (See [Figure 10-56](#).)

Figure 10-56 WLAN Pre-authentication ACL

The specific ACL is configured under Security > Access Control Lists (See [Figure 10-57](#) and [Figure 10-58](#).)

Figure 10-57 WLC Access Control Lists


Note

If a pre-authentication ACL is used in conjunction with the web auth policy, it must include a rule to permit DNS requests; otherwise, the client is unable to resolve and connect to a destination host/URL that is otherwise allowed by the ACL.

Figure 10-58 Pre-Auth ACL Example

Security

AAA

RADIUS

Authentication

Accounting

TACACS+

LDAP

Local Net Users

MAC Filtering

Disabled Clients

User Login Policies

AP Policies

Local EAP

Priority Order

Access Control Lists

Access Control Lists

CPU Access Control Lists

Access Control Lists > Edit

< Back

Add New Rule

General

Access List NameCisco Open Garden

Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port	DSCP	Direction
1	Permit	10.20.31.0 / 255.255.255.0	0.0.0.0 / 0.0.0.0	UDP	Any	DNS	Any	Any
2	Permit	0.0.0.0 / 0.0.0.0	10.20.31.0 / 255.255.255.0	UDP	DNS	Any	Any	Any
3	Permit	10.20.31.0 / 255.255.255.0	171.71.181.19 / 255.255.255.255	TCP	Any	HTTP	Any	Any
4	Permit	171.71.181.19 / 255.255.255.255	10.20.31.0 / 255.255.255.0	TCP	HTTP	Any	Any	Any

221882

Anchor Controller DHCP Configuration

If the anchor controller is going to manage DHCP services for the guest access WLAN, proceed with the steps below.


Note

The anchor controller cannot be used to manage DHCP services if guest N+1 redundancy is being implemented, because there is no mechanism to synchronize address leases for a single guest VLAN/subnet across two or more WLCs.

Adding a New DHCP Scope to the Anchor Controller


- Step 1

Click the **Controller** tab.
- Step 2

In the left pane, click **Internal DHCP Server**.

Step 3 Click **New**. (See [Figure 10-59](#).)

Figure 10-59 Adding a New DHCP Scope

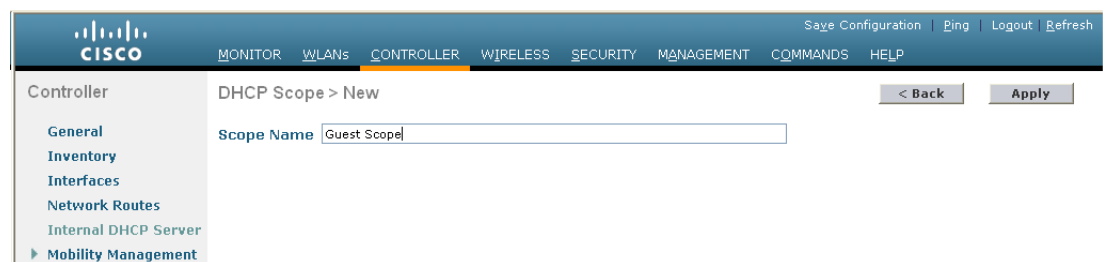


221859

Defining a Scope Name

Step 4 Define a name for the scope and click **Apply**. (See [Figure 10-60](#).)

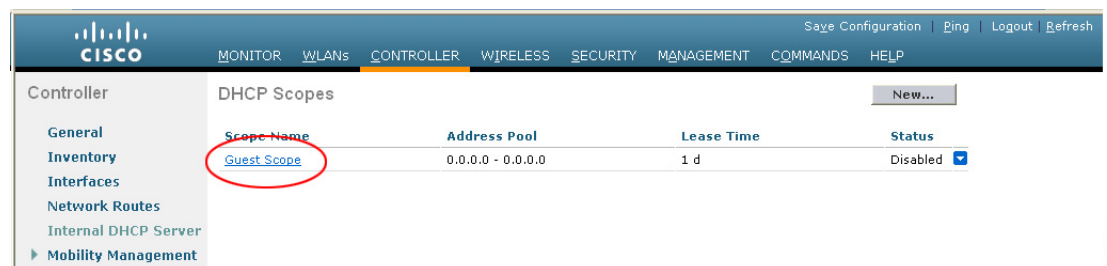
Figure 10-60 Defining a Scope Name



221859

Step 5 Click **Scope Name** to edit. (See [Figure 10-61](#).)

Figure 10-61 Editing DHCP Scope



221860

Defining Scope Properties

Step 6 Define the following minimum information:

- Pool start and stop
- Network
- Mask
- Default routers
- DNS servers

Step 7 For Status, select **Enabled** and click **Apply**. (See [Figure 10-62](#).)

Figure 10-62 Configuring and Enabling Scope Properties

The screenshot shows the Cisco Unified Wireless Network Controller configuration page for a DHCP Scope. The page is titled "DHCP Scope > Edit" and includes a "Back" button and an "Apply" button. The configuration fields are as follows:

Field	Value
Scope Name	Guest Scope
Pool Start Address	10.20.31.100
Pool End Address	10.20.31.200
Network	10.20.31.0
Netmask	255.255.255.0
Lease Time (seconds)	86400
Default Routers	10.20.31.1, 0.0.0.0, 0.0.0.0
DNS Domain Name	
DNS Servers	171.68.226.120, 171.70.168.183, 0.0.0.0
Netbios Name Servers	0.0.0.0, 0.0.0.0, 0.0.0.0
Status	Enabled

221861

External Radius Authentication

As described in [Guest User Authentication](#), an external RADIUS server can be used to authenticate guest users in place of creating and storing guest credentials locally on the anchor controller. If this method is used, the lobby admin features described in [Guest Account Management](#) cannot be used. It is assumed that some other guest management system will be used in conjunction with the external RADIUS server.

To configure a guest WLAN to use an external RADIUS server, perform the following configuration steps on the anchor controller.

Adding a RADIUS Server

Step 1 Click the **Security** tab.

A summary screen is displayed. (See [Figure 10-63](#).)

Figure 10-63 Summary Screen

Security

RADIUS Authentication Servers

Call Station ID Type: IP Address

Credentials Caching: ☐

Use AES Key Wrap: ☐ (Designed for FIPS customers and requires a key wrap compliant RADIUS server)

Network User	Management	Server Index	Server Address	Port	IPsec	Admin Status
<input type="checkbox"/>	<input checked="" type="checkbox"/>	1	10.20.30.16	1812	Disabled	Enabled
<input type="checkbox"/>	<input checked="" type="checkbox"/>	2	10.20.30.15	1812	Disabled	Enabled

Step 2 Click **New**.

The screen shown in [Figure 10-64](#) appears.

Figure 10-64 Defining RADIUS Server Settings

Security

RADIUS Authentication Servers > New

Server Index (Priority): 3

Server IP Address: 10.20.30.17

Shared Secret Format: ASCII

Shared Secret:

Confirm Shared Secret:

Key Wrap: ☐ (Designed for FIPS customers and requires a key wrap compliant RADIUS server)

Port Number: 1812

Server Status: Enabled

Support for RFC 3576: Enabled

Retransmit Timeout: 2 seconds

Network User: ☐ Enable

Management: ☐ Enable

IPsec: ☐ Enable

Step 3 To define RADIUS server settings, configure the IP address, shared secret, and authentication port number as defined on the RADIUS server.

If the Network User check box is cleared, the RADIUS server is used only for user authentication when it is specifically selected under the RADIUS setting of a given WLAN. Otherwise, if the Network User check box is checked, the server is used globally for all user authentications based on its server priority.

Step 4 Click **Apply**.

The summary screen shown in [Figure 10-65](#) shows the newly-added server.

Figure 10-65 Summary Screen

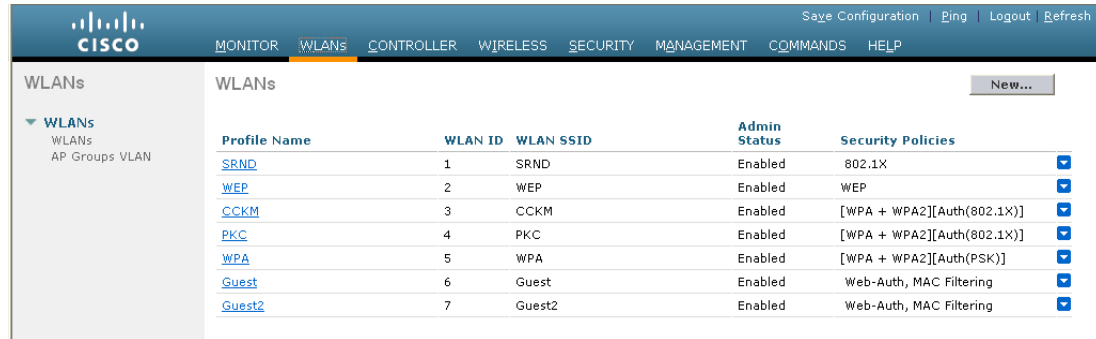
The screenshot shows the Cisco Unified Wireless Network Guest Access Services configuration interface. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh. The main menu on the left lists various security and management options, with 'RADIUS' selected under 'AAA'. The main content area is titled 'RADIUS Authentication Servers' and includes an 'Apply' button and a 'New...' link. Below this, there are configuration options for 'Call Station ID Type' (set to 'IP Address'), 'Credentials Caching' (unchecked), and 'Use AES Key Wrap' (unchecked, with a note that it is designed for FIPS customers and requires a key wrap compliant RADIUS server). A table lists three RADIUS servers with their respective configurations.

Network User	Management	Server Index	Server Address	Port	IPSec	Admin Status
<input type="checkbox"/>	<input checked="" type="checkbox"/>	1	10.20.30.16	1812	Disabled	Enabled <input checked="" type="checkbox"/>
<input type="checkbox"/>	<input checked="" type="checkbox"/>	2	10.20.30.15	1812	Disabled	Enabled <input checked="" type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	3	10.20.30.17	1812	Disabled	Enabled <input checked="" type="checkbox"/>

221909

- Step 5** To select a RADIUS server, click the **WLANS** tab.
The screen shown in [Figure 10-66](#) appears.

Figure 10-66 *WLANS Tab*

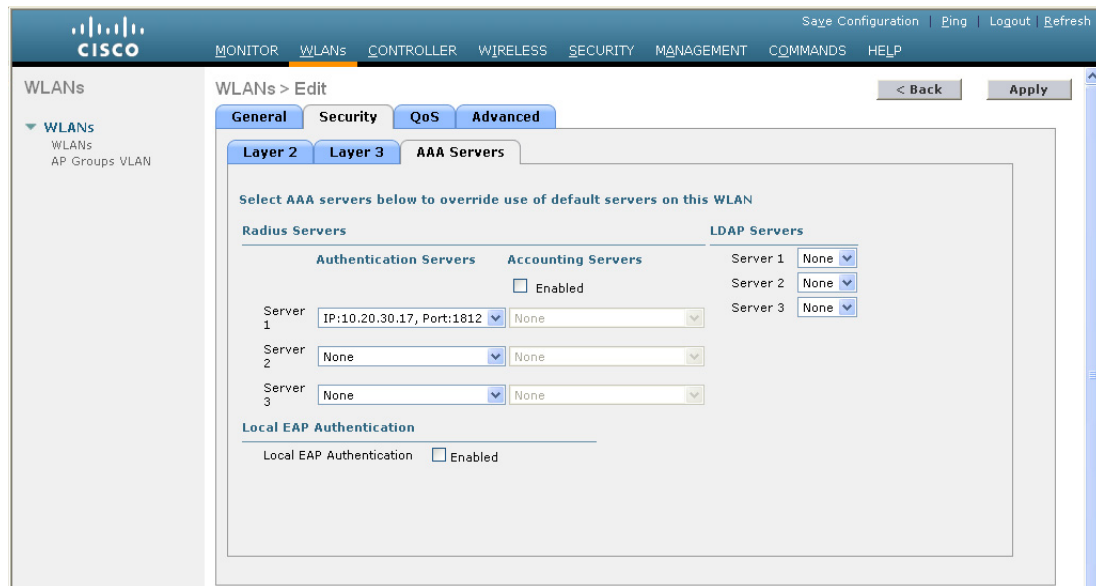


Profile Name	WLAN ID	WLAN SSID	Admin Status	Security Policies
SRND	1	SRND	Enabled	802.1X
WEP	2	WEP	Enabled	WEP
CCKM	3	CCKM	Enabled	[WPA + WPA2][Auth(802.1X)]
PKC	4	PKC	Enabled	[WPA + WPA2][Auth(802.1X)]
WPA	5	WPA	Enabled	[WPA + WPA2][Auth(PSK)]
Guest	6	Guest	Enabled	Web-Auth, MAC Filtering
Guest2	7	Guest2	Enabled	Web-Auth, MAC Filtering

221910

- Step 6** Find the guest WLAN and click on its **Profile Name**.
The guest WLAN configuration screen is displayed, as shown in [Figure 10-67](#).

Figure 10-67 *Guest WLAN Configuration Screen*



WLANs > Edit

General Security QoS Advanced

Layer 2 Layer 3 AAA Servers

Select AAA servers below to override use of default servers on this WLAN

Radius Servers		LDAP Servers	
Authentication Servers	Accounting Servers	Server 1	Server 2
Server 1: IP:10.20.30.17, Port:1812	None	None	None
Server 2: None	None	None	None
Server 3: None	None	None	None

Local EAP Authentication

Local EAP Authentication ☐ Enabled

221911

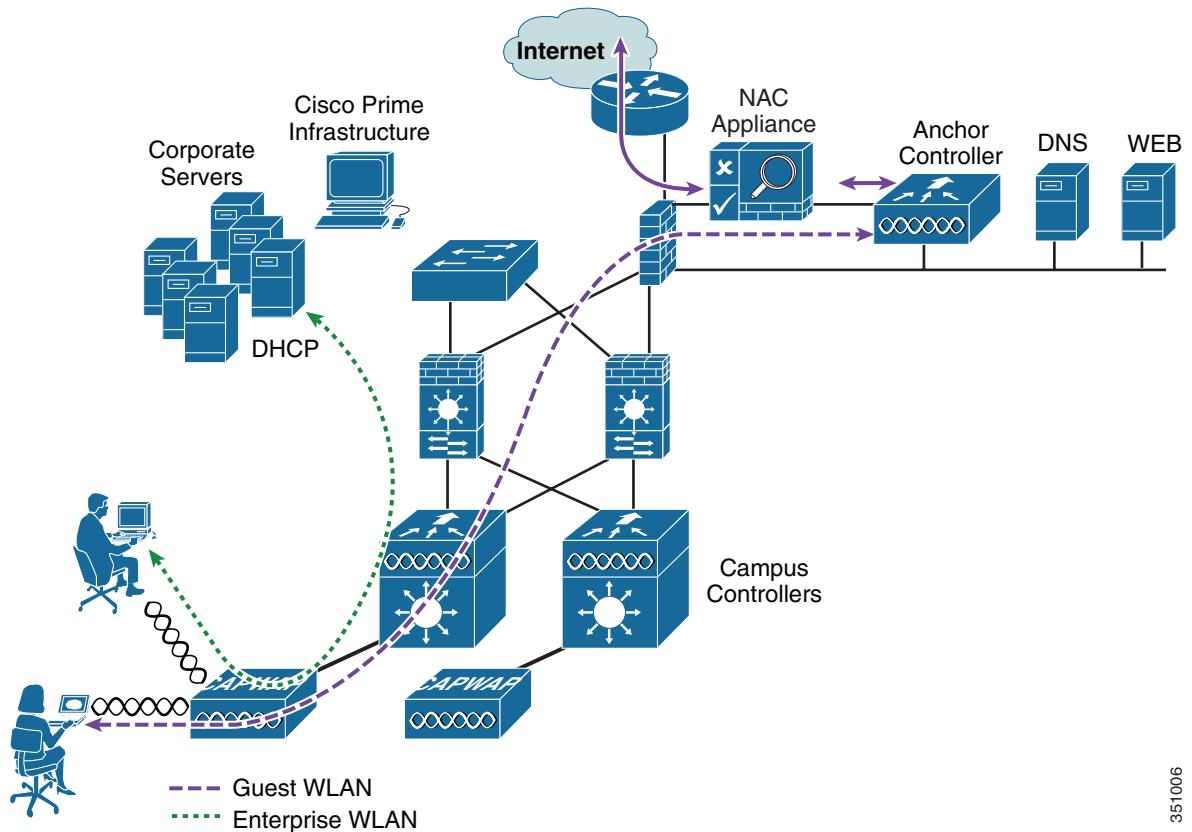
- Step 7** Select **AAA Servers** under the WLAN Security tab
- Step 8** Select the RADIUS server to be used for web authentication from the pull-down selection list under Authentication Servers.

External Access Control

The centralized guest access topology described in this chapter can be integrated with an external access control platform such the Cisco NAC Appliance.

In this scenario, an enterprise might have already deployed an access control platform in their Internet DMZ to support wired guest access services (see [Figure 10-68](#)).

Figure 10-68 *Wireless Guest Access with External Access Control*



As shown in [Figure 10-68](#), the wireless guest access topology remains the same except that the guest VLAN interface on the anchor controller, instead of connecting to a firewall or border router, connects to an inside interface on an access control platform such as the Cisco NAC Appliance.

In this scenario, the NAC Appliance is responsible for redirection, web authentication, and subsequent access to the Internet. The campus and anchor controllers are used only to tunnel guest WLAN traffic across the enterprise into the DMZ, where the NAC appliance or some other platform is used to control guest access.

Configuration of the guest WLAN, campus, and anchor controllers is the same as described in the previous examples.

The only exception is that Layer 3 web policy is not enabled under the guest WLAN security settings (see [Figure 10-69](#) and [Figure 10-70](#)).

Figure 10-69 Guest WLAN Layer 3 Security Policy

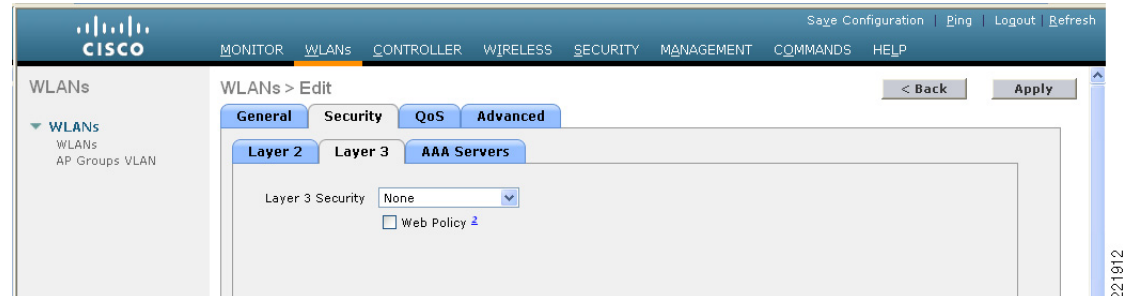
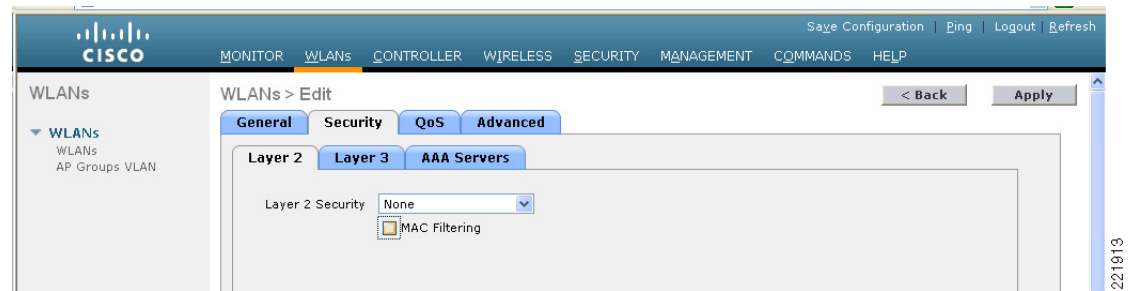


Figure 10-70 Guest WLAN L2 Security Settings



The configurations above establishes a WLAN with no security policies. Guest traffic passes through the anchor controller to the inside or untrusted interface of the Cisco NAC Appliance, where it is blocked until the user has authenticated.

DHCP can be hosted locally on the controller or externally via the NAC Appliance or dedicated server.

It is beyond the scope of this chapter to address Cisco NAC Appliance or other external access control platform specific configurations. See the specific platform documentation for additional configuration guidelines.

Verifying Guest Access Functionality

The guest access service is working correctly if a user:

- Can associate to the guest WLAN
- Receives an IP address via DHCP
- Opens their browser and is redirected to the web authentication page
- Enters their credentials and connects to the Internet (or other authorized upstream services)

Troubleshooting Guest Access

The following verifications and troubleshooting tasks assume the following:

- The solution is using the web authentication functionality resident in the anchor controller(s).
- User credentials are created and stored locally on the anchor controller(s).

Before attempting to troubleshoot the various symptoms below, at the very least you should be able to ping from the campus (foreign) controller to the anchor controller(s). If not, verify routing.

Next, you should be able to perform the following advanced pings. These can only be performed via the serial console interfaces of the controllers:

- **mping** *neighbor WLC ip*

This pings the neighbor controller through the CAPWAP control channel.

- **eping** *neighbor WLC ip*

This pings the neighbor controller through the CAPWAP data channel.

If a standard ICMP ping goes through, but mpings do not, ensure that the default mobility group name of each WLC is the same, and ensure that the IP, MAC, and mobility group name of each WLC is entered in the mobility members list of every WLC.

If pings and mpings are successful, but epings are not, check the network to make sure that IP protocol 97 (Ethernet-over-IP) is not being blocked.

User Cannot Associate to the Guest WLAN

- Verify that the guest WLAN is enabled on the anchor controller and all foreign controllers that support the guest WLAN
- Verify that the guest WLAN SSID is being broadcast.
- Verify client adapter/software configuration.

User Does Not Obtain an IP Address via DHCP

- Verify that WLAN configuration settings are identical on the anchor and foreign controllers (except for WLAN interface and mobility anchors; see [Guest WLAN Configuration on the Anchor WLC, page 10-26](#))
- Verify that the guest WLAN is enabled on the anchor WLC(s)
- Check for a proper DHCP server address under the guest VLAN interface settings on the anchor controller(s)
 - If using an external DHCP server, the IP address should be that of the external server.
 - Verify reachability to the external DHCP server from the anchor controller.
 - If using the anchor controller for DHCP services, the DHCP server IP address should be the management IP address of the controller.
 - Verify that a DHCP scope has been configured and enabled on the controller.
 - Verify that the network mask of the DHCP scope is consistent with the mask on the guest VLAN interface.
 - Verify that the DHCP scope does not overlap with any addresses assigned to the network infrastructure.

User is Not Redirected to Web Auth Page

The following assumes the user is able to associate to the guest WLAN and obtain an IP address:

- Verify that valid DNS servers are being assigned to the client via DHCP.
- Ensure that the DNS servers are reachable from the anchor controller.
- Verify that the URL being opened in the web browser is resolvable.
- Verify that the URL being opened in the web browser is connecting to HTTP port 80.



Note The internal web auth server does not redirect incoming requests on ports other than 80 and one other user defined port number (see [User Redirection, page 10-8](#)).

User Cannot Authenticate

- Verify that user credentials are active on the anchor controller(s).
Guest credentials typically have a lifetime associated with them. If the credentials have expired, they do not appear under the Security > Local Net Users list on the anchor controller. Use Cisco Prime Infrastructure to re-apply the user template or re-create user credentials locally on the controller. See [Guest Management Using the Management System](#) and [Guest Credentials Management](#).
- Verify user password.

User Cannot Connect to Internet or Upstream Service

- Verify routing to and from the anchor controller from the firewall or border router connecting to the anchor controller(s)
- Verify NAT configuration on firewall or Internet border router (if applicable)

System Monitoring

Following are some monitoring commands that might be helpful in troubleshooting.

Anchor Controller

From the serial console port:

```
Cisco Controller) >show client summary
Number of Clients..... 1
MAC Address      AP Name      Status      WLAN  Auth  Protocol  Port
-----
00:40:96:ac:5f:f8 10.15.9.19   Associated   3     Yes  Mobile    1
```

Note that the protocol is mobile. The Auth field reflects the actual status of the user. If the user has passed web auth, the field displays YES. If not, the field shows NO.

Also notice the AP name. This is the management IP address of the foreign controller (originating controller).

From the summary information, use the client MAC to show more detail:

```
(Cisco Controller) >show client detail 00:40:96:ac:5f:f8
Client MAC Address..... 00:40:96:ac:5f:f8
Client Username ..... romaxam
AP MAC Address..... 00:00:00:00:00:00
Client State..... Associated
Wireless LAN Id..... 3
BSSID..... 00:00:00:00:00:02
Channel..... N/A
IP Address..... 10.20.31.100
Association Id..... 0
Authentication Algorithm..... Open System
Reason Code..... 0
Status Code..... 0
Session Timeout..... 86316
Client CCX version..... No CCX support
Mirroring..... Disabled
QoS Level..... Silver
Diff Serv Code Point (DSCP)..... disabled
802.1P Priority Tag..... disabled
WMM Support..... Disabled
Mobility State..... Export Anchor
Mobility Foreign IP Address..... 10.15.9.19
Mobility Move Count..... 1
Security Policy Completed..... Yes
Policy Manager State..... RUN
Policy Manager Rule Created..... Yes
NPU Fast Fast Notified..... Yes
Policy Type..... N/A
Encryption Cipher..... None
Management Frame Protection..... No
EAP Type..... Unknown
Interface..... wlan-user
VLAN..... 31
Client Capabilities:
    CF Pollable..... Not implemented
    CF Poll Request..... Not implemented
    Short Preamble..... Not implemented
    PBCC..... Not implemented
    Channel Agility..... Not implemented
    Listen Interval..... 0
Client Statistics:
    Number of Bytes Received..... 0
    Number of Bytes Sent..... 0
    Number of Packets Received..... 0
    Number of Packets Sent..... 0
    Number of Policy Errors..... 0
    Radio Signal Strength Indicator..... Unavailable
    Signal to Noise Ratio..... Unavailable
Nearby AP Statistics:
    TxExcessiveRetries: 0
    TxRetries: 0
    RtsSuccessCnt: 0
    RtsFailCnt: 0
    TxFiltered: 0
    TxRateProfile: [0,0,0,0,0,0,0,0,0,0,0,0]
```

The same information can be obtained through the web configuration and management interface of the controller under Clients > Detail. (See [Figure 10-71](#).)

Figure 10-71 Anchor WLC Monitor > Client Detail

The screenshot shows the Cisco WLC Monitor > Client Detail page. The page is divided into several sections:

- Client Properties:**
 - MAC Address: 00:40:96:ac:5f:f8
 - IP Address: 10.20.31.100
 - Client Type: Regular
 - User Name: romaxam
 - Port Number: 1
 - Interface: wlan-user
 - VLAN ID: 31
 - CCX Version: Not Supported
 - E2E Version: Not Supported
 - Mobility Role: Export Anchor
 - Mobility Peer IP Address: 10.15.9.19
 - Policy Manager State: RUN
 - Mirror Mode:
 - Management Frame Protection: No
- AP Properties:**
 - AP Address: Unknown
 - AP Name: 10.15.9.19
 - AP Type: Mobile
 - WLAN Profile: Guest2
 - Status: Associated
 - Association ID: 0
 - 802.11 Authentication: Open System
 - Reason Code: 0
 - Status Code: 0
 - CF Pollable: Not Implemented
 - CF Poll Request: Not Implemented
 - Short Preamble: Not Implemented
 - PBCC: Not Implemented
 - Channel Agility: Not Implemented
 - Timeout: 0
 - WEP State: WEP Disable
- Security Information:**
 - Security Policy Completed: Yes
 - Policy Type: N/A
 - Encryption Cipher: None
 - EAP Type: N/A
- Quality of Service Properties:**
 - WMM State: Disabled

Campus (Foreign) Controller

From the serial console port:

```
(WiSM-slot3-1) >show client summary
Number of Clients..... 2
MAC Address      AP Name      Status      WLAN  Auth  Protocol  Port
-----
00:40:96:ac:5f:f8  AP3_.18e5.7fdc  Associated   1     Yes  802.11g   29
```

Note that the protocol field is 802.11g, whereas the protocol field on the anchor controller for the same client is mobile. The campus (foreign) controller always shows the user as authenticated and the AP name reflects the actual AP to which the client is associated.

Additional details can be obtained using the following:

```
(WiSM-slot3-1) >show client detail 00:40:96:ac:5f:f8
Client MAC Address..... 00:40:96:ac:5f:f8
Client Username ..... N/A
AP MAC Address..... 00:17:df:35:86:50
Client State..... Associated
Wireless LAN Id..... 1
BSSID..... 00:17:df:35:86:50
Channel..... 11
IP Address..... Unknown
Association Id..... 1
Authentication Algorithm..... Open System
Reason Code..... 0
Status Code..... 0
Session Timeout..... 0
Client CCX version..... No CCX support
Mirroring..... Disabled
QoS Level..... Silver
Diff Serv Code Point (DSCP)..... disabled
```

```

802.1P Priority Tag..... disabled
WMM Support..... Disabled
Mobility State..... Export Foreign
Mobility Anchor IP Address..... 10.15.9.13
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..... management
VLAN..... 9
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..... 308244
    Number of Bytes Sent..... 700059
    Number of Packets Received..... 2527
    Number of Packets Sent..... 1035
    Number of Policy Errors..... 0
    Radio Signal Strength Indicator..... -75 dBm
    Signal to Noise Ratio..... 25 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]
    AP3_.18e5.7fdc(slot 0) .....
antenna0: 37 seconds ago -73 dBm..... antenna1: 4294510568 seconds ago
-128 dBm

```

The same information can be obtained through the controller web configuration and management interface under Clients > Detail (see [Figure 10-72](#)).

Figure 10-72 Foreign WLC Monitor > Client Detail

The screenshot displays the Cisco WLC Monitor > Client Detail page. The page is divided into several sections:

- Client Properties:**
 - MAC Address: 00:40:96:ac:5f:f8
 - IP Address: 0.0.0.0
 - Client Type: Regular
 - User Name:
 - Port Number: 29
 - Interface: management
 - VLAN ID: 9
 - CCX Version: Not Supported
 - E2E Version: Not Supported
 - Mobility Role: Export Foreign
 - Mobility Peer IP Address: 10.15.9.13
 - Policy Manager State: RUN
 - Mirror Mode:
 - Management Frame Protection: No
- AP Properties:**
 - AP Address: 00:17:df:35:86:50
 - AP Name: AP3_18e5.7fdc
 - AP Type: 802.11g
 - WLAN Profile: Guest2
 - Status: Associated
 - Association ID: 1
 - 802.11 Authentication: Open System
 - Reason Code: 0
 - Status Code: 0
 - CF Pollable: Not Implemented
 - CF Poll Request: Not Implemented
 - Short Preamble: Implemented
 - PBCC: Not Implemented
 - Channel Agility: Not Implemented
 - Timeout: 0
 - WEP State: WEP Disable
- Security Information:**
 - Security Policy Completed: Yes
 - Policy Type: N/A
 - Encryption Cipher: None
 - EAP Type: N/A
- Quality of Service Properties:**
 - WMM State: Disabled

The page also includes a left sidebar with navigation options: Summary, Statistics, CDP, and Wireless. The top navigation bar includes links for Save Configuration, Ping, Logout, and Refresh. The bottom right corner shows the date 22/9/19.

Debug Commands

Additional debug commands that might be used from the serial console include the following:

```
debug mac addr <client mac address>
debug mobility handoff enable
debug mobility directory enable
debug dhcp packet enable
debug pem state enable
debug pem events enable
debug dot11 mobile enable
debug dot11 state enable
```




Cisco Mobility Services Engine

Introduction

This chapter provides configuration and deployment guidelines if you want add the Cisco Mobility Services Engine (MSE) and run Context Aware Services in a Cisco Unified Wireless Network. The purpose of this chapter is to:

- Explain the various elements and framework of the Cisco Mobility Solution
- Provide general deployment guidelines



Note

This chapter does not provide configuration details for the MSE and associated components. This information is provided in other documents. See [Related Information](#) for a list of documents about the configuration and design of Context Aware Mobility Services. Adaptive wIPS configuration is also not covered in this design guide.

Background Information

The Cisco MSE provides the ability to track the physical location of Network Devices, both wired and wireless, using wireless LAN controllers (WLCs) and Cisco Aironet CAPWAP APs. This solution allows you to track any Wi-Fi device, including clients, active RFID tags, and rogue clients and APs. It was designed with the following requirements:

- **Manageability**—Cisco Prime Infrastructure is used to administer and monitor the MSE. Moreover, the MSE integrates directly into the wireless LAN architecture, which provides one unified network to manage instead of multiple disparate wireless networks.
- **Scalability**—The Cisco MSE series can simultaneously track 25,000 elements in CAS and 5,000 APs in wIPS. The CPI can manage multiple Mobility Services Engines for greater scalability. The wireless LAN controller (WLC), CPI, and MSE are implemented through separate devices to deliver greater scalability and performance.
- **Security**—The WLC, CPI, and MSE provide robust secure interfaces and secure protocols to access data. The MSE records historical location information that can be used for audit trails and regulatory compliance.
- **Open and standards based**—The MSE has a SOAP/XML API that can be accessed by external systems and applications that can leverage location information from the MSE.

- Easy deployment of business applications—The MSE can be integrated with new business applications such as asset tracking, inventory management, location-based security, or automated workflow management.

Overview

Context Aware Service (CAS) provides the capability for a Wi-Fi 802.11a/b/g/n network to determine the location of a person or object with an active Wi-Fi device, such as a wireless client or active RFID tag and/or associated data that can be passed by the end point through the wireless infrastructure to an upstream client. When a Cisco MSE is added to a Cisco Unified Wireless Network with an appropriately licensed version of CPI, it assumes responsibility for several important tasks:

- Execution of positioning algorithms
- Maintenance of calibration information
- Trigger and dispatch of location notifications
- Process of statistics and historical location
- Depository for geographical information, maps, and all wireless devices

Cisco Prime Infrastructure is the management system that interfaces with the MSE and serves as the user interface (UI) for the services that the MSE provides. Although it is possible to access the MSE directly through SSH or a console session for maintenance and diagnostic purposes, all operator and user interaction with the MSE is typically performed through CPI (for management) or a third-party location client application.

Terminology

With the Cisco *Centralized WLAN Architecture* (functional architecture of Cisco Unified Wireless Networks) and Context Aware Location Services, administrators can determine the location of any 802.11-based device as well as the specific type or status of each device. Clients (associated, probing, and so forth), rogue APs, rogue clients, and active tags can all be identified and located by the system. This information is made available through the API within seconds of an event occurrence and can be retained by the MSE database for historical lookup and security audits.

Mobility Services Engine

MSE supports a suite of mobility services applications. Designed as an open platform, the MSE supports mobility services software in a modular fashion with various configuration options based on network topology and the types of services required. Cisco supports existent and future applications that include:

- Context Aware Services—These applications capture and integrate into business processes detailed contextual information about such things as location, temperature, availability, and applications used. Context Aware applications feature a wide range of location options that include real-time location, presence detection, chokepoint visibility, and telemetry. Support for enhanced received signal strength indication (RSSI) and Time Difference of Arrival (TDOA) technology delivers greater scale accuracy and performance for a broad range of environments.

Context Aware software consists of two major components:

- Context Aware Engine for Clients: RSSI, the Cisco location engine, is used to track Wi-Fi clients, rogue clients, rogue APs, and wired clients.

- Context Aware Engine for Tags: The partner (AeroScout) location engine (both RSSI and TDOA) is used to track Wi-Fi active RFID tag.

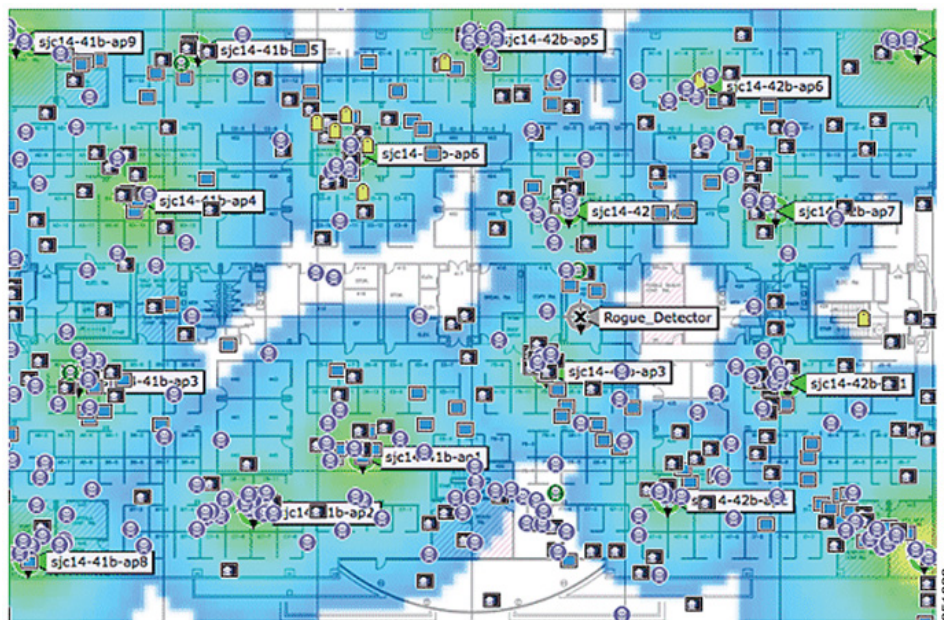
Third-party software is supported through the MSE API.

- Adaptive Wireless Intrusion Prevention System (wIPS)—wIPS software provides visibility and comprehensive threat prevention for the mobility network through monitoring, alerts, classifying, and remediation of wireless and wired network vulnerabilities.
- Network Mobility Services Protocol—Cisco-defined protocol that is used for secure communication between the WLC and MSE.
- Cisco Prime Infrastructure—Wireless network management system developed and supported by Cisco Systems. Includes these capabilities:
 - WLAN configuration
 - WLAN performance monitoring
 - Reporting (real-time and historical)
 - Graphical view of network (wireless LAN controllers, APs, clients and tags)
- Wireless LAN Controller (WLC)—The centralized Cisco Unified Wireless Network architecture configuration and control device. This controller allows the entire WLAN to operate as an intelligent network that uses wireless as the access medium to support advanced services, unlike legacy 802.11 WLAN infrastructures that are built from autonomous, discrete APs. The Cisco Unified Wireless Network simplifies operational management by collapsing large numbers of managed endpoints—autonomous APs—into a single managed system comprised of one or more WLCs and their corresponding, joined APs.

In the Cisco Unified Wireless Network architecture, APs are *lightweight*, which means that they cannot act independently of a WLC. APs are typically *zero-touch* deployed, and no individual configuration of APs is required. The APs learn the IP address of one or more WLCs through a controller discovery algorithm and then establish a trust relationship with a controller through a join process. Once the trust relationship is established, the WLC pushes firmware to the AP, if necessary, and a run-time configuration. APs do not store a configuration locally.

- Clients—All devices associated with controller-based, lightweight APs on a wireless network.
- Rogue Access Point—Any AP that is determined not to be part of the wireless LAN mobility group that detected it. This consists of all non-system APs within RF range of a lightweight APs, which includes those on the wired network or those on another wired network (such as an AP of a neighbor). Because all lightweight APs use a hash as part of the beacon frame with a special key, even spoofed infrastructure APs are identified as rogue APs rather than mistaken to be legitimate APs flagged in CPI as spoof APAPs.
- Rogue Clients—All devices that are associated to rogue APs.
- Active RFID Tags—Wi-Fi device that can be detected and located on a Wi-Fi network. There is wide variety of Wi-Fi compatible tags available in the market. Tags offer a range of features that include telemetry, such as motion and environmental data such as temperature and humidity, call buttons, indoor and outdoor operation, intrinsically safe versions, and flexible mounting options.

The MSE provides the ability to track up to 25,000 devices (tags, clients, and rogue clients/APs). [Figure 11-1](#) is an example of a floor map as shown in the CPI, and displays tags, clients, rogue clients and rogue APs. The floor map illustrates the scale and variety of classes of devices that can be tracked by the MSE. CPI provides the capability to define search parameters to display only in a subset of devices. For example, a biomedical user might want to see only infusion pumps and EKG machines named with friendly identifiers rather than rogue devices or devices with cryptic MAC or IP addresses.

Figure 11-1 CPI Floor Map with Tracked Devices

Map Legend:

- Client—Blue square monitor
- Tag—Yellow vertical rectangle
- Rogue AP—Circle with skull-and-crossbones (red = malicious, green = friendly, gray = unclassified)
- Rogue Client—Blue square monitor with skull-and-crossbones

Technology Background Information

There are two technologies that are used to track Wi-Fi devices with the Cisco Mobility Solution:

- RSSI—Received signal strength indication
- TDOA—Time difference of arrival

For details on these technologies, see *Wi-Fi Location-Based Services Design Guide*, at:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/WiFiLBS-DG.html>



Note

This reference is to an older release but the information is still valid.

RSSI

RSSI is the measured power of a received radio signal. The packets transmitted by any wireless device are received at multiple APs (provided that those APs listen on the channel on which the frame was transmitted). The APs forward these packets to the wireless LAN controller along with the correspondent RSSI information measured at the AP. The wireless LAN controller aggregates this information on a per device basis from different APs. This data is forwarded to the MSE through NMSP. The Context Aware Services that reside on the MSE use the RSSI data received from one or more WLCs

to determine the location of a wireless device. RSSI is usually preferred for indoor or low ceiling environments, which can result in reflection of the signals. Unlike TDOA, RSSI does not require exact time synchronization amongst APs. With the measured RSSI values from different APs, the probability of the location of a device is calculated at different points on the floor. Based on this probability, the location is returned as the estimated location.

Time Difference on Arrival

The time difference on arrival (TDOA) mechanism is the preferred method to determine device location when you track tags in outdoor and outdoor-like environments, such as indoor, high-ceiling environments. With TDOA, the location of a WLAN device is determined based on the time of arrival of the signal that it transmits as seen by three or more time-synchronized Wi-Fi TDOA receivers. The time of arrival data is collected and reported to the Context Aware Engine for Tags that reside on the MSE, which computes the time differences of arrival between multiple pairs of Wi-Fi TDOA receivers. The time required for a given message to be received by different Wi-Fi TDOA receivers is proportional to the length of the transmission path between the mobile transmitting device and each TDOA receiver. This mechanism of calculation device location requires time synchronization between the Wi-Fi TDOA receivers.

In order to compute a position accurately, this method requires a set of at least three Wi-Fi TDOA receivers. The distance between Wi-Fi TDOA receivers is relatively larger than the distance between APs that are required for indoor RSSI positioning. As with RSSI positioning, this method relies on unidirectional communication (tag transmitting notification frame, no association required).

Refer to the *Context Aware Service Software Configuration Guide*, at:

http://www.cisco.com/en/US/docs/wireless/mse/3350/7.0/CAS/configuration/guide/CAS_70.html

Active RFID Tags

Cisco Compatible Extensions (CCX)-compliant active RFID tags are detected on a Wi-Fi network based on tag notification frames that are sent by the tag and received by an 802.11 AP. The tag notification frame rate can be programmed based on the specific use case scenario. Typically, tags are configured to transmit tag notification frames every 3-5 minutes to optimize frequent location updates and battery life. The call button feature provides the ability to trigger events based on push button on the tag. This enables advanced functionality, such as emergency reporting or parts replenishment. Some tags provide more than one call button. The second call button can be programmed for additional functionality. Tags can store pre-programmed messages that can be received by the wireless network infrastructure. A battery is used to power active tags, which provides up to four years of battery life. Battery life is dependent upon a number of tag configuration parameters that includes the frequency of tag notification frame transmission and repetition rate. Tags can report on their battery level and alert when low. Tags can also have a built-in motion sensor to transmit tag notification frames upon movement. This helps to conserve battery life when the tag is stationary; configure the tags to transmit less frequently when they do not move.

There is another category of tags that add advanced sensor technology to accurately monitor the condition of an asset, such as its ambient temperature, in addition to other location and status information. These sensor tags use standard Wi-Fi networks to transport the asset location and sensor data and do not require dedicated or proprietary sensor networks.

Wi-Fi RFID tags that are compliant with the CCX for Wi-Fi Tags specification can optionally pass tag telemetry information to the location-aware Cisco UWN as part of their tag message payload. Telemetry information is received by APs and collected by the WLCs. At MSE startup, the MSE subscribes for all the service in which it is interested, such as the measurements for tags. The WLC continues to send the MSE notifications at the end of each aggregation cycle.

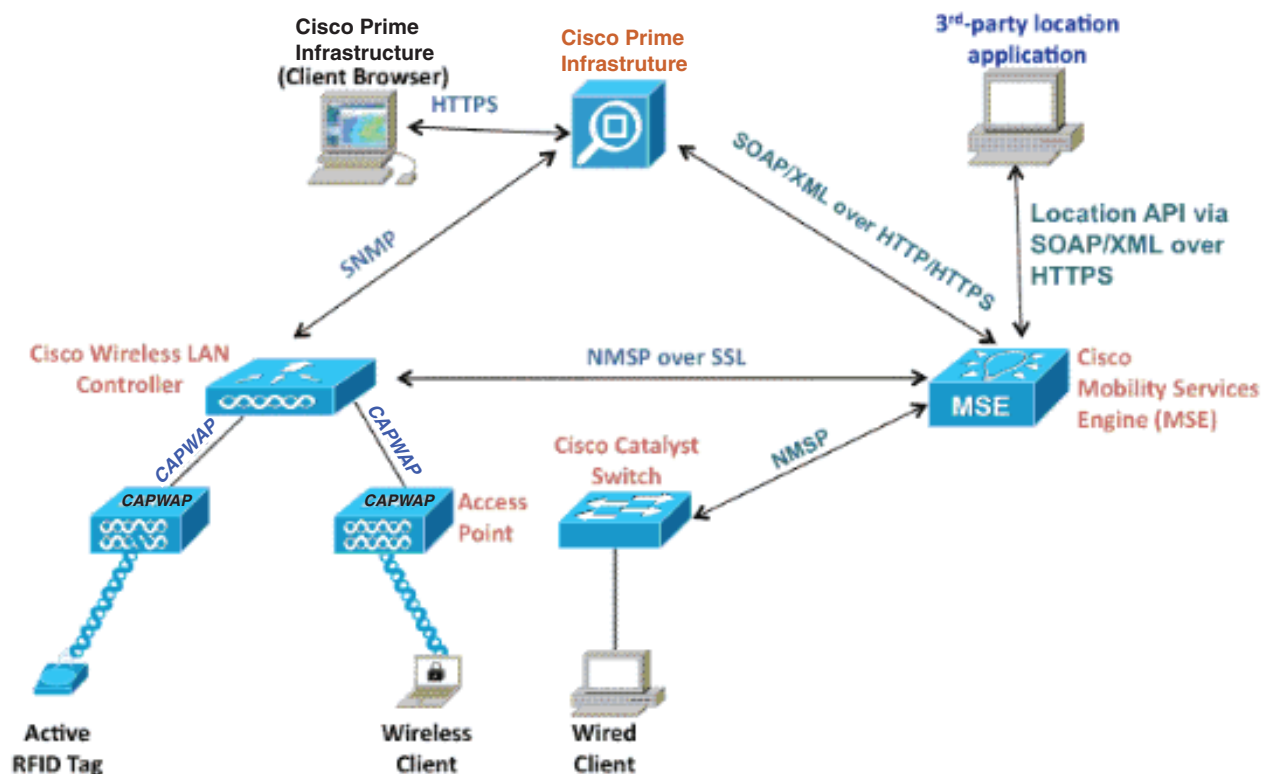
Telemetry information is transmitted from a CCX-compatible tag and is received by one or more APs and/or location receivers, that is, Wi-Fi TDOA receivers, which, in turn, pass the telemetry information to their respective registered WLAN controllers. If the tags are configured to send multiple frame copies (or bursts) per channel, the controller eliminates any duplicate tag telemetry and passes the distilled telemetry values to the MSE. The database in the MSE is updated with the new telemetry information and makes it available to location clients through the MSE SOAP/XML API.

In the case of a tag that passes telemetry value, NMSP is designed to efficiently transport telemetry values from multiple tags in a similar fashion. Telemetry traffic from multiple tags is aggregated by the WLC with each NMSP endpoint capable of performing NMSP frame fragmentation and reassembly if required. All tag data can be included in the northbound notifications, which includes telemetry, call buttons, chokepoint encounters, and so forth.

System Architecture

The MSE integrates with the Cisco centralized wireless LAN architecture as shown in [Figure 11-2](#). The MSE sits out of the data path of the wireless LAN and receives data from the WLC through NMSP. Cisco Prime Infrastructure is used to configure the MSE. Once configured, the MSE is self-contained.

Figure 11-2 System Architecture



When you deploy the Context Aware solution, consideration must be given to the type of devices tracked and the maximum device count. You can track any of the five device types (Wi-Fi clients, active RFID tags, rogue clients, rogue APs, or wired clients) to be configured individually or for simultaneous tracking.

One MSE can be managed by only one CPI, that is, a single MSE cannot be managed by multiple CPI instances, but a single CPI can manage multiple MSEs. When the number of devices to be managed exceeds the capacity of a single MSE, you need to deploy multiple, independent MSEs. The ability to deploy multiple MSEs for scaling applies to all services currently supported on MSE. The maximum number of devices that can be tracked by one Cisco MSE 3355 is 25,000 devices (combination of Wi-Fi clients, active RFID tags, rogue clients, rogue APs, and wired clients) as part of Context Aware Service. The older model Cisco MSE 3310 can track up to 2,000 devices, while the Cisco MSE 3350 can track up to 18,000 devices. When the number of devices to be managed exceeds the capacity of a single MSE box, multiple, independent MSE appliances need to be deployed. This can require MSEs on specific controllers, especially on large campuses where roaming of clients or assets can cross different physical buildings or domains. In this instance, controllers can communicate with a maximum of 10 MSE appliances.

Cisco lightweight APs operate in a unique dual mode that detect devices both on the channels where they service clients and also on all other channels if they periodically background scan while still provide data access to their wireless clients. The gathered raw location data is then forwarded from each AP to its associated WLC through the LWAPP or standards-based CAPWAP protocol. Data is transported between the WLC and the MSE through a secure NMSP connection.

Cisco Prime Infrastructure is used to manage and configure the MSE, and it can also become the visual front-end of the MSE to display Wi-Fi devices that are tracked. All device (wired and wireless) details and specific historical location information can be accessed with the MSE northbound API. CPI uses this interface to visualize location information, as well as view and configure Context Aware parameters.

The Cisco Mobility Solution consists of two location engines with a single unified API. The location engines are:

- Context Aware Engine for Clients (Cisco engine)—Can be used for both Clients and Tags
- Context Aware Engine for Tags (partner engine)—AeroScout-based TAG solution

The **Context Aware Engine for Clients** is an RSSI-based solution that is ideal for tracking Wi-Fi client devices in indoor spaces (for example, offices, hospitals, or other low-ceiling environments). This engine ships by default on all Cisco MSE servers. The tracking licenses of the Context Aware Engine for Clients can be shared between Clients and Tags.

The **Context Aware Engine for Tags** has the ability to use both an RSSI and TDOA-based engine and is intended to be used when you track Wi-Fi devices in indoor, low-ceiling (RSSI), indoor high-ceiling (TDOA), and outdoor (TDOA) environments. This engine is also installed by default on all MSE platforms and is license enabled. The tracking licenses of the Context Aware Engine for Tags can be used only for Tags. For client tracking you need to purchase the following additional components:

- Tag tracking license for the MSE with appropriate Tag count (TDOA or RSSI)
- Wi-Fi TDOA location receivers (if and when required)
- LR license for each Wi-Fi TDOA receiver

When a Cisco MSE is added to a Cisco Unified Wireless Network, the MSE assumes responsibility for the following important tasks:

- Execution of positioning algorithms
- Maintenance of calibration information
- Triggering and dispatch of location notifications
- Processing of statistics and historical location

Cisco Prime Infrastructure is the management platform for the MSE servers and is the user interface (UI) for the services that the MSE provides. The MSE is accessed directly through SSH or a console session for maintenance and diagnostic purposes. All operator and user interaction with the MSE is usually through CPI. The integration of a Cisco MSE into a Cisco Unified Wireless Network architecture immediately enables improvements to base-level location capabilities, which include:

- **Scalability**—If you add a Cisco MSE, it increases the scalability of the Cisco Unified Wireless Network from on-demand tracking of a single device at a time to a maximum tracking capacity of up to 25,000 simultaneous devices (WLAN clients, RFID tags, rogue APs, and rogue clients) per MSE. For deployments that require support of greater numbers of devices, additional MSE appliances can be deployed and managed under one or more CPI servers.
- **Historical and statistics trending**—The MSE records and maintains historical location and statistics information for clients and tags. This information is available for viewing through CPI or with third-party location clients. This historical information can be used for location trending, asset loss investigation, RF capacity management, and facilitation of network problem resolution. Historical parameters are configured in Cisco Prime Infrastructure.

Related Information

The following references provide additional information about the Cisco Mobility Services Engine.

- *Cisco 3355 Mobility Services Getting Started Guide*

http://www.cisco.com/en/US/docs/wireless/mse/3355/user/guide/mse_qsgmain.html

- *Cisco 3350 Mobility Services Getting Started Guide*

http://www.cisco.com/en/US/docs/wireless/mse/3350/quick/guide/mse_qsgmain.html

- *Cisco 3310 Mobility Services Engine Getting Started Guide*

http://www.cisco.com/en/US/docs/wireless/mse/3310/quick/guide/MSE3310_GSG.html

- *Cisco Mobility Services Engine - Context Aware Mobility Solution Deployment Guide*

http://www.cisco.com/en/US/products/ps9742/products_tech_note09186a00809d1529.shtml

- *Cisco Context Aware Service Configuration Guide*

http://www.cisco.com/en/US/docs/wireless/mse/3350/6.0/CAS/configuration/guide/msecg_ch7_CAS.html

- *Cisco 3300 Series Mobility Services Engine Licensing and Ordering Guide*

http://www.cisco.com/en/US/docs/wireless/mse/3350/7.3/CAS_Configuration_Guide/Guide/msecg_Ovview.html

- *Wi-Fi Location-Based Services 4.1 Design Guide* (old release but information is still useful)

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/WiFiLBS-DG.html>

- Mobility Groups FAQ

http://www.cisco.com/en/US/products/ps6366/products_qanda_item09186a00809a30cc.shtml



A

AAA	Authentication, Authorization, and Accounting.
ACS	Cisco Access Control Server.
AES	Advanced Encryption Standard.
AP	Access point.

B

BSSID	Basic service set identifier.
--------------	-------------------------------

C

CAM	Clean Access Manager.
CCMP	Counter Mode with Cipher Block Chaining Message Authentication Code Protocol.
CCX	Cisco Compatible Extensions.
CKIP	Cisco Key Integrity Protocol.
CMIC	Cisco Message Integrity Check.
CSA	Cisco Security Agent.
CSSC	Cisco Secure Services Client. Cisco Key Integrity Protocol (CKIP) and Cisco Message Integrity Check (CMIC).

D

DoS	Denial of service.
------------	--------------------

E

EAP	Extensible Authentication Protocol.
EAP-FAST	EAP-Flexible Authentication via Secured Tunnel.

EAP-TLS	EAP-Transport Layer Security.
EIRP	Effective Isotropic Radiated Power.
ESSID	Extended service set identifier, commonly referred to as an SSID.

F

FWSM	Firewall Services Module.
-------------	---------------------------

I

IDS	Intrusion detection system.
IPS	Intrusion prevention system.

L

LAP	LWAPP Access Point.
LBS	Location-based service
LWAPP	Lightweight Access Point Protocol.

M

MAP	Mesh AP
MFP	Management frame protection.
MIC	Message integrity check.

N

NAC	Network Admission Control.
------------	----------------------------

O

OFDM	Orthogonal Frequency Division Multiplexing.
-------------	---

P

PEAP GTC	Protected EAP Generic Token Card.
PEAP MSCHAP	Protected EAP Microsoft Challenge Handshake Authentication Protocol.
PKI	Public Key Infrastructure.

R

RADIUS	Remote Authentication Dial-In User Service.
RF	Radio frequency.
RFID	Radio frequency. Radio-frequency identification.
RLDP	Rogue Location Discovery Protocol.
RSSI	Received signal strength indication.

S

SNR	Signal-to-noise ratio.
SSID	IEEE Extended Service Set Identifier.
SSO	Single sign-on.
SVI	Switched virtual interfaces.

T

TKIP	Temporal Key Integrity Protocol.
TLS	Transport Layer Security.

W

WCS	Wireless Control System.
WEP	Wired Equivalent Privacy.
Wi-Fi	Wi-Fi is the brand of the Wi-Fi Alliance, which certifies interoperability of products and services based on IEEE 802.11 technology.
WiSM	Wireless Services Module.

WLAN	Wireless LAN.
WLC	Wireless LAN Controller.
WLCM	Wireless LAN Controller Module.
WLSM	Wireless LAN Services Module.
WMM	Wi-Fi Multimedia
WPA	Wi-Fi Protected Access.