


# Cisco Connected Mobile Experiences (CMX) CVD

February 10, 2015



Building Architectures to Solve Business Problems



## About the Authors



Zeb Hallock

### Zeb Hallock, Technical Marketing Engineer, Systems Development Unit, Cisco Systems

Zeb Hallock is in the Enterprise Systems Engineering group of Cisco, focusing on digital media systems. He is also pursuing creation and development of future-based collaboration systems, holding two patents in the field. He has been with Cisco for 10 years working on enterprise system testing, system design and testing of H.323 based video conferencing, and network infrastructure. He has also been a specialist working on Cisco Unified IP Contact Center Cisco Unified MeetingPlace. Before Cisco he worked as a consultant designing and implementing local and wide area networks.



Suyog Deshpande

### Suyog Deshpande, Technical Marketing Engineer, Systems Development Unit, Cisco Systems

Suyog Deshpande is a Technical Marketing Engineer within Systems Development Unit (SDU) at Cisco Systems, focussing on converged wired and wireless access products. Suyog has been with Cisco for 9 years with previous experience in the Wireless Networking Business Unit, where he focussed on RF systems and converged access products. Before Cisco, Suyog held various positions designing and deploying large scale wireless networks and working with spectrum intelligence and analyzer products.



Roland Saville

### Roland Saville, Technical Leader, Systems Development Unit, Cisco Systems

Roland Saville is a technical leader for the Systems Development Unit (SDU) at Cisco, focused on developing best-practice design guides for enterprise network deployments. He has more than 15 years of experience at Cisco as a systems engineer, consulting systems engineer, technical marketing engineer, and technical leader. During that time, he has focused on a wide range of technology areas including the integration of voice and video onto network infrastructures, network security, wireless LAN networking, RFID, and energy management. He has also spent time focusing on the retail market segment. Prior to Cisco, he spent eight years as a communications analyst for Chevron Corporation. Saville holds a bachelor of science degree in electrical engineering from the University of Idaho and a master of business administration degree from Santa Clara University. He co-authored the book "Cisco TelePresence Fundamentals" and has eight U.S. patents.

### Stephenie Chastain, Senior Technical Manager, Systems Development Unit, Cisco Systems

Stephenie Chastain is a Senior Technical Manager in the Systems Development Unit (SDU) at Cisco. Her focus is on business networks in the area of IP network design, primarily in education. She has been with Cisco for more than 12 years and has broad experience in product management,



Stephenie Chastain

customer support and IP networks including design and implementation of large service provider broadband access networks and enterprise network designs. Her educational background includes applied physics from University of Georgia and electrical engineering from Georgia Institute of Technology.



# About Cisco Validated Design (CVD) Program

---

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit <http://www.cisco.com/go/designzone>.

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.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned 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. (1110R).

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Connected Mobile Experiences (CMX) CVD

© 2014 Cisco Systems, Inc. All rights reserved.



## About Cisco Validated Design (CVD) Program v

---

### PART 1

## CMX Introduction

### Preface 3

---

### CHAPTER 1

## Connected Mobile Experiences Solution Overview 1-1

Introduction 1-1

Connected Mobile Experiences (CMX) 1-2

CMX and the Wireless Infrastructure 1-2

Role of the Mobile User in the CMX Network 1-3

Role of the Organization in the CMX Network 1-3

Concerns for CMX and Mobile Users 1-3

CMX Solution Advantages 1-4

---

### PART 2

## CMX Design Overview

---

### CHAPTER 2

## Summary of CMX Design Overview 2-1

---

### CHAPTER 3

## CMX Solution Components 3-1

Wireless Infrastructure 3-1

    Cisco Aironet Access Points 3-1

    Cisco Wireless LAN Controllers (WLCs) 3-3

    Cisco Mobility Services Engine (MSEs) 3-3

    Cisco Prime Infrastructure 3-4

    Cisco Context Aware Service (CAS) 3-4

    Probe Request RSSI versus FastLocate 3-7

Connected Mobile Experiences Services 3-9

    CMX Location Analytics 3-9

    Dashboard Tab 3-11

    CMX Analytics Tab 3-11

    CMX Reports 3-12

    Differences in CMX Reports, Dashboard, and Analysis 3-13

    CMX Presence Analytics 3-13

    CMX Visitor Connect 3-14

---

### CHAPTER 4

## CMX Deployment Models 4-1

    Overview 4-1

- Deployment Topologies 4-1
- WAN Bandwidth Utilization 4-3
- MSE Scalability 4-7
- Campus and Branch Designs 4-9
  - Single Campus (or Large Branch) Deployment Model 4-10
  - Small Branch Deployment Model 4-12
  - B2C Guest Access for CMX Visitor Connect 4-13

---

**CHAPTER 5**

**CMX Security Considerations 5-1**

- Traffic Isolation for CMX Visitor Connect 5-1
- Role-Based Access Control on the MSE 5-4
- Role-Based Access Control for the CMX Connect & Engage Service 5-6

---

**CHAPTER 6**

**CMX Additional Considerations 6-1**

- Currency of Location Information 6-1
- Apple iOS Version 8 Mobile Devices 6-2
- Android Mobile Devices 6-3
- 2.4 GHz vs. 5 GHz Mobile Devices 6-4
- FastLocate Deployment Restrictions 6-4

---

**PART 3**

**CMX Use Case Stories**

---

**CHAPTER 7**

**CMX Use Case Stories 7-1**

- CMX Location Analytics Use Case Stories 7-1
- CMX Visitor Connect Use Case Story 7-3
- CMX Presence Analytics Use Case Story 7-3

---

**PART 4**

**CMX Radio Frequency and Location Based Design**

---

**CHAPTER 8**

**Summary of CMX Radio Frequency and Location Based Design 8-1**

---

**CHAPTER 9**

**Radio Operating Frequencies and Data Rates 9-1**

- Radio Frequency Bands 9-1
- Regulatory Domains 9-1
- Operating Frequencies 9-2
- 802.11 Modulation Techniques 9-2
- Direct-Sequence Spread Spectrum 9-3

Frequency-Hopping Spread Spectrum	9-3
Orthogonal Frequency Division Multiplexing	9-3
2.4 GHz Operating Frequencies and Data Rates	9-3
5 GHz Operating Frequencies and Data Rates	9-4
802.11ac	9-6

**CHAPTER 10****Radio Frequency Fundamentals 10-1**

Power Level	10-1
Effective Isotropic Radiated Power	10-3
Path Loss	10-4
Receive Signal Strength Indicator—RSSI	10-4
Signal to Noise Ratio—SNR Ratio	10-5
Signal Attenuation	10-5
Example Use Case	10-6

**CHAPTER 11****Antenna Fundamentals 11-1**

Antenna Gain	11-1
Antenna Types	11-2
Omnidirectional Antenna	11-2
Directional Antennas	11-2
Multipath Distortion	11-3
Diversity Antenna Systems and Multipath Distortion	11-4
Antenna Orientation and Access Point Placement	11-4
Defining Individual Access Point Heights	11-5

**CHAPTER 12****802.11 Fundamentals 12-1**

Beacons	12-1
802.11 Join Process—Association	12-2

**CHAPTER 13****Location Fundamentals 13-1**

Probe RSSI	13-2
Location Latency	13-3
FastLocate	13-5
Location Accuracy and Currency	13-6

**CHAPTER 14**

**Pre-Deployment Radio Frequency Site Survey 14-1**

- Pre-deployment RF Site Survey 14-1
- Physical Site Survey 14-1
  - Location Assessment 14-1
- Business Needs of WLAN 14-2
- Constraints on Deployment 14-3
- Budgeting 14-3
- Existing 802.11 Surveys 14-4
- Use Case Example 14-4

**CHAPTER 15**

**Access Point Placement and Separation 15-1**

- Access Point Separation 15-1
- AP Placement 15-5
- Proper Access Point Placement 15-7
- Improper Access Point Placement 15-9
- Getting Around Placement 15-10
- Recommend Access Point Placement 15-12

**CHAPTER 16**

**Predictive Radio Frequency Planning 16-1**

- Cisco Prime Infrastructure RF Planning Tool 16-2
- Ekahau RF Planning 16-8

**CHAPTER 17**

**Multi-Floor Deployments 17-1**

- Limited Flexibility for Placing APs 17-1
- Inter-floor Interference Issues 17-1
- AP Deployment Guidelines to Mitigate Inter-floor Issues 17-2
- Multi-Floor RF Site Survey 17-5
  - Hospitals 17-6
  - Warehouses 17-7
  - Manufacturing Facility 17-7

**CHAPTER 18**

**Capacity Planning and High Density 18-1**

- Access Point Density 18-1
- High Density Deployment 18-2
- Establish and Validate a Per-Connection Bandwidth Requirement 18-4
- Calculate the Aggregate Throughput Required for the Coverage Area 18-5

- 802.11 and Scalability—How Much Bandwidth Will a Cell Provide? 18-6
- Other High Density Considerations 18-7

**CHAPTER 19****Location Voice and Data Co-Existence 19-1**

- Minimum Desired Cell Signal Level Threshold 19-1
- Signal to Noise Ratio (SNR) 19-1
- Data Rate 19-2
- Cell-to-Cell Overlap 19-2

**CHAPTER 20****Post-Deployment Radio Frequency Tuning 20-1**

- Radio Resource Management 20-1
- Transmit Power Control 20-2
- Overriding the TPC Algorithm with Minimum and Maximum Transmit Power Settings 20-3
- Dynamic Channel Assignment 20-3
- Coverage Hole Detection and Correction 20-4
- Benefits of RRM 20-5
- CleanAir 20-5
  - Role of the Cisco Wireless LAN Controller in a Cisco CleanAir System 20-6
  - Interference Types that Cisco CleanAir Can Detect 20-6
  - Persistent Devices 20-7
  - Persistent Devices Detection 20-7
  - Persistent Devices Propagation 20-8
  - Detecting Interferers by an Access Point 20-8
- Post-Deployment RF Tuning 20-8
  - Location Assessment 20-8
  - Business Needs of WLAN 20-9
  - Constraints on Deployment 20-9
  - Existing 802.11 Surveys 20-9

**CHAPTER 21****Best Practices Checklist 21-1****PART 5****CMX Configuring the Infrastructure****CHAPTER 22****Summary of CMX Configuring the Infrastructure 22-1****CHAPTER 23****Configuring Cisco Wireless LAN Controllers 23-1**

- WLC Visitor Connect Configuration 23-1



Configuring the ACL for CMX Visitor Connect 23-1  
 Configuring the WLAN for Visitor Connect 23-2  
 Configuring FastLocate 23-3

**CHAPTER 24**

**Configuring Cisco Prime Infrastructure 24-1**

Installing Cisco Prime Infrastructure 24-1  
 Installing the Cisco Mobility Services Engine 24-1  
 Adding Wireless LAN Controllers to Cisco Prime Infrastructure 24-1  
 Configuring Maps within Cisco Prime Infrastructure 24-5  
     Adding Floor Areas to a Campus Building or a Standalone Building 24-5  
     Adding APs on Maps 24-9  
         Adding Access Points to a Floor Area 24-9  
         Defining Coverage Area 24-12  
     Monitoring Geo-Location 24-14  
         Adding a GPS Marker to a Floor Map 24-14  
         Editing a GPS Marker 24-15  
         Deleting a GPS Marker Present on a Floor 24-15  
     Inclusion and Exclusion Areas on a Floor 24-15  
         Defining an Inclusion Region on a Floor 24-16  
         Defining an Exclusion Region on a Floor 24-16  
 WebGL Requirements 24-17  
 Adding Mobility Services Engine 24-18  
 Synchronizing Controller and Network Designs 24-21

**CHAPTER 25**

**Configuring the Mobility Services Engine for CMX 25-1**

Verifying CMX Settings 25-2  
 Configuring Role-Based Access Control (RBAC) on the MSE 25-5

**CHAPTER 26**

**Configuring CMX Analytics 26-1**

Logging In 26-1  
 Configuring CMX Presence Analytics 26-2  
     Threshold Settings 26-3  
     Importing Access Points 26-5  
     Adding Presence Sites 26-5  
 Configuring CMX Location Analytics 26-7  
     Configuring the CMX Analytics Dashboard 26-8  
         Adding a New Page 26-8  
         Modifying or Deleting an Existing Page 26-13

	Customizing CMX Analysis	26-15
	Zone Analysis	26-21
	Alternative Paths Analysis	26-25
	Heat Maps	26-27
	Typical Locations	26-28
	Customizing CMX Reports	26-29
	Conversion Percentage Report	26-29
	Daily Visitors and Dwell Time Report	26-31
	Detected versus Connected Devices Report	26-32
	Hourly Visitors and Dwell Time Report	26-32
	Movement between Zones Report	26-33
	Repeat Visitors Report	26-34
<b>CHAPTER 27</b>	<b>Configuring CMX Visitor Connect</b>	<b>27-1</b>
	Configuring CMX Visitor Connect with Splash Pages and Social Connectors	27-1
	Configuring Facebook App for Visitor Connect	27-8
	Visitor Policy	27-10
	Server Settings	27-11
	Configuring RBAC on CMX Connect & Engage	27-12
<b>PART 6</b>	<b>CMX Appendices</b>	
<b>APPENDIX A</b>	<b>CMX Software Versions</b>	<b>A-1</b>
<b>APPENDIX B</b>	<b>CMX System Release Notes</b>	<b>B-1</b>
	MSE 8.0 Role-Based Access Control	B-1
<b>APPENDIX C</b>	<b>802.11 Data Rates</b>	<b>C-1</b>
	IEEE 802.11a/n/ac	C-1
	IEEE 802.11b/g/n	C-2
	Maximum Power Levels and Antenna Gains	C-3
	IEEE 802.11a	C-3
	IEEE 802.11b	C-3
<b>APPENDIX D</b>	<b>CMX Use Case Example—Upgrade VoWLAN Ready Network to Location/CMX Ready</b>	<b>D-1</b>
<b>APPENDIX E</b>	<b>CMX Troubleshooting</b>	<b>E-1</b>
	MSE and WLC Communication Problems	E-1

[Aspect Ratio Issues while Creating Maps](#) E-2

[Coverage Zones Cannot Be Renamed](#) E-2



## **PART 1**

### **CMX Introduction**





# Preface

---

**September 4, 2014**

This document is a Cisco Validated Design (CVD) for Cisco Connected Mobile Experience (CMX) Solutions. It presents system-level requirements, recommendations, guidelines, and best practices for detecting, connecting, and engaging mobile users within your venue and leveraging your Wi-Fi network to fit your business needs. As Cisco continues to develop and enhance the technologies required to implement a CMX solution, this CVD will evolve and be updated to provide the latest guidelines, recommendations, and best practices for designing and deploying a CMX solution.

## How to Use this Document

This document is organized into five main parts after the initial [Chapter 1, “Connected Mobile Experiences Solution Overview.”](#)

## CMX Design Overview

The chapters in this part of the document describe the main components of Cisco CMX solution and explain how these components work together to form a complete end-to-end solution:

- [Chapter 3, “CMX Solution Components”](#)—Highlights the wireless (Wi-Fi) network infrastructure necessary for providing location services and CMX services within this design guide.
- [Chapter 4, “CMX Deployment Models”](#)—Introduces high-level models for the deployment of infrastructure components necessary for location services and CMX. Considerations around bandwidth utilization and scalability of the MSE are discussed.
- [Chapter 5, “CMX Security Considerations”](#)—Focuses on traffic isolation for guest wireless access as part of CMX Visitor Connect and also discusses the CMX Connect and Engage service.
- [Chapter 6, “CMX Additional Considerations”](#)—Highlights the additional considerations when deploying a CMX Solution such as fast location information, Apple IOS 8 devices, and considerations around 2.4 and 5GHZ frequency bands when deploying location services and CMX services.

## CMX Use Cases

The chapter in this part of the document describes the CMX use case examples tested and verified within this design guide. A chapter summarizing the use cases verified as well as Video On Demand (VoD)s of each of the use cases showcase the CMX solution components in several real world scenarios.

- [Chapter 7, “CMX Use Case Stories”](#)—Introduces several use cases that can be met through the deployment of CMX services. Each is designed to highlight the application of CMX services to address a realistic business scenario. The first two use cases involve the use of CMX Location Analytics with a large-sized retail scenario to analyze customer behavior to provide better service. The final two use cases involve the use of the CMX Visitor Connect service, as well as the use of CMX Presence Analytics within small-sized retail scenarios, to provide customer Wi-Fi access as well as analyze customer behavior to provide better service.

## RF and Location Based Design

The chapters in this part of the document describe various services in addition to the use cases described in the previous section:

- [Chapter 9, “Radio Operating Frequencies and Data Rates”](#)—Discusses RF operating frequencies that are used for WLAN Networks. 802.11 a/b/n/ac modulation techniques and the role of TPC and DCA in a RF network are discussed.
- [Chapter 10, “Radio Frequency Fundamentals”](#)—Discusses RF fundamentals that must be understood before deploying a Wireless LAN network that is location and CMX ready. The chapter explains various RF concepts such as spectrum bands, power level, signal strength, RSSI, etc.
- [Chapter 11, “Antenna Fundamentals”](#)—Discusses antennas, which are a fundamental part of any WLAN deployment, and how selecting the right type of antenna for deployment greatly enhances both coverage and location readiness.
- [Chapter 12, “802.11 Fundamentals”](#)—Discusses 802.11 fundamentals, namely the role of beacons, probe requests, and probe responses.
- [Chapter 13, “Location Fundamentals”](#)—Discusses location fundamentals, including definition of a location ready point, location currency, location accuracy, and location latency. We also discuss two methods of obtaining location from a client, i.e., the Probe RSSI method and the FastLocate method.
- [Chapter 14, “Pre-Deployment Radio Frequency Site Survey”](#)—Discusses the pre-deployment RF site survey. A good Cisco WLAN deployment is dependent on a good RF design, including doing a thorough site survey of the location, determining the best location for access points, making the right channel plans, planning for AP capacity, and performing a regular post deployment RF site survey
- [Chapter 15, “Access Point Placement and Separation”](#)—Discusses AP placement and AP capacity planning. Core concepts regarding the distance between APs in a network and its impact on location data and voice are discussed. Additionally designing for capacity is also discussed in the chapter.
- [Chapter 16, “Predictive Radio Frequency Planning”](#)—Discusses predictive RF planning that should be under taken after a pre-deployment RF Site Survey is completed. Two tools to perform RF planning are discussed, namely the Cisco Prime Infrastructure RF Planner tool and the Ekahau Site Survey tool.
- [Chapter 17, “Multi-Floor Deployments”](#)—Discusses challenges in deployments that involve multiple floors. Recommendations on what to keep in mind while designing for RF network are also discussed.



- [Chapter 18, “Capacity Planning and High Density”](#)—Discusses planning a network with capacity and need in mind. Today’s WLAN needs are heavily dependent on mobile devices and applications. Capacity planning involves looking at application needs and designing a network around them, while High Density networks may be required when too many clients are expected to connect in a location.
- [Chapter 19, “Location Voice and Data Co-Existence”](#)—Discusses the pertinent characteristics of voice and data designs only as they relate to co-existence with the location tracking capabilities of the Cisco UWN.
- [Chapter 20, “Post-Deployment Radio Frequency Tuning”](#)—Discusses post-deployment RF tuning that should be done regularly on the deployment and includes using RRM for channel planning, CleanAir to mitigate RF interference, and a regular post site survey assessment to ensure that optimum RF health is maintained.
- [Chapter 21, “Best Practices Checklist”](#)—Discusses the best practices check list while deploying a CMX solution.

## Configuring the Infrastructure

The chapters in this part of the document describe the network infrastructure design and configuration foundations to deploy a CMX solution in a customer environment:

- [Chapter 23, “Configuring Cisco Wireless LAN Controllers”](#)—Highlights the configuration of the Business to Consumer (B2C) guest WLAN necessary for providing guest wireless connectivity leveraging CMX Visitor Connect. In addition, information regarding how to enable presence on the WLC is provided.
- [Chapter 24, “Configuring Cisco Prime Infrastructure”](#)—Highlights the configuration of Cisco Prime Infrastructure for map integration and management of the CMX solution.
- [Chapter 25, “Configuring the Mobility Services Engine for CMX”](#)—Highlights the configuration of the MSE to collect and report the location-based information and analytics.
- [Chapter 26, “Configuring CMX Analytics”](#)—Highlights the configuration options for the three main functional areas of CMX Analytics—Dashboard, Analytics and Reports.
- [Chapter 27, “Configuring CMX Visitor Connect”](#)—Highlights the configuration required to enable CMX Visitor Connect on the MSE.

## Appendices

The appendices contain useful information that is not covered in the main chapters of this CVD:

- [Appendix A, “CMX Software Versions”](#)—Provides the software versions and devices leveraged in this design guide.
- [Appendix B, “CMX System Release Notes”](#)—Provides important information you should be aware of when designing and implementing this release of the CMX CVD.
- [Appendix C, “802.11 Data Rates”](#)—Lists the data rates for 802.11an/ac rates in 5GHz and 802.11bgn in 2.4 GHz.
- [Appendix D, “CMX Use Case Example—Upgrade VoWLAN Ready Network to Location/CMX Ready”](#)—Provides a use case to transition your already existing WLAN network to a location ready network to deploy CMX.

- [Appendix E, “CMX Troubleshooting”](#)—Provides several troubleshooting techniques when deploying this release of the CMX CVD.

## For Experienced Users

Readers who are familiar with previous versions of this CVD or who are experienced at designing a CMX solution can use this document as a reference source. Rather than reading every page or every chapter, this document has been broken into modules that can be easily searched for a particular topic. Updates to the topics in this CVD will be published periodically.

## For New Users

This document is long and contains an extensive amount of complex technical information. It can seem intimidating, particularly if you are a first time reader of this document or do not have much experience with a CMX solution.

To orient yourself to the document, we recommend you begin with [Chapter 2, “Summary of CMX Design Overview,”](#) which provides an overview of the major components required to deploy a CMX solution and typical use cases. From this section, you can then determine if you need particular design guidance around the infrastructure, the uses cases, or a set operation.

## Where to Find Additional Information

Because the document covers a wide spectrum of Cisco Network Infrastructure, Security, and Mobility products and possible solution designs, it cannot provide all the details of individual products, features, or configurations. For that type of detailed information, refer to the specific product documentation available at: <http://www.cisco.com>.

This document provides general guidance on how to design your own CMX solution. Cisco has developed, tested, and documented specific solutions for certain applications and has made those solutions available for customers to copy and deploy. They are part of the Cisco Validated Design program described and documented at: <http://www.cisco.com/go/designzone>.

## Revision History

This document may be updated at any time without notice. You can obtain the latest version of this document online at: [TBD](#).

Visit this website periodically and check for documentation updates by comparing the revision date of your copy with the revision date of the online document.

[Table 1](#) lists the revision history for this document.

**Table 1**      **Revision History**

Revision Date	Comments
September 3, 2014	Initial version of this CMX CVD.

# Command Syntax Conventions

Table 2 describes the syntax used with the commands in this document.

**Table 2**      **Command Syntax Guide**

<b>Convention</b>	<b>Description</b>
<b>boldface</b>	Commands and keywords.
<i>italic</i>	Command input that is supplied by you.
[ ]	Keywords or arguments that appear within square brackets are optional.
{ x   x   x }	A choice of keywords (represented by x) appears in braces separated by vertical bars. You must select one.
^ or Ctrl	Represent the key labeled <i>Control</i> . For example, when you read ^D or <i>Ctrl-D</i> , you should hold down the Control key while you press the D key.
screen font	Examples of information displayed on the screen.
<b>boldface screen font</b>	Examples of information that you must enter.
< >	Nonprinting characters, such as passwords, appear in angled brackets.
[ ]	Default responses to system prompts appear in square brackets.





# Connected Mobile Experiences Solution Overview

---

September 4, 2014

## Introduction

The emergence of ubiquitous wireless networks and the explosion of mobile devices means that nearly everyone has access to the Internet and can be contacted through a communications or data network. Mobile devices are no longer used solely for workplace activities that expand productivity and reshape work habits, but are now used as an instant source of information for users. Users leverage their mobile devices to discover, compare, share, and communicate information about products and services. With the increase of users leveraging their mobile devices, organizations have a new way to deliver innovative user services and enhance the customer experience by leveraging their wireless networks in their venues.

In addition to ubiquitous access to Wi-Fi networks and the growth of the smart phone and smart tablets, the industry is also seeing the growth of the mobile application. With the widespread acceptance of mobile applications, users worldwide take it for granted that they can access information anywhere and at any time. Access to Wi-Fi is expected and the market is seeing the emergence of a new phenomenon: the mobile connected user. Today's mobile connected users are bringing their smart devices into a venue and using them to look up prices, find information, and post to social media. This behavior has created opportunities for organizations to utilize their existing IT Wi-Fi network to connect with their customers to increase both loyalty and revenue.

Built on Cisco's WLAN infrastructure, Cisco's Connected Mobile Experiences (CMX) allows enterprises and service providers to deliver customized, location-based mobile services that not only provide a timely, personalized mobile device experience, but also enable organizations to better understand their users through onsite, online, and social analytics. Location-based services allow mobile users to receive useful information or capabilities based on their location within a venue. In turn, organizations acquire information about their mobile users that allows them to provide better services as well as track the success of their engagement strategies.

Organizations that recognize this new class of mobile users and the benefits the CMX solution offers can increase their revenue by providing personalized and relevant information based on where a user is located in the venue. With CMX solutions, organizations are able to:

- Build customer intimacy, loyalty, and retention.
- Elevate venue operations with intelligent product placement, appropriate staffing, and improved floor layouts.

- Transition Wi-Fi from an IT expenditure into a profit center through third-party mobile advertising opportunities and mobile-influenced sales.

## Connected Mobile Experiences (CMX)

To build relationships with their users and ultimately increase revenue, organizations must:

- Engage users—Organizations must find new ways to reach their users that go beyond passive marketing campaigns. Today's mobile user wants relevant content, information, and services delivered to their mobile devices based on their location and personal preferences.
- Improve the user experience—Organizations must find new ways to provide an unprecedented user experience to increase customer satisfaction and loyalty, such as listing services inside venues and delivering customized information directly to a mobile device based on the user's location.
- Understand user behavior—Attracting a user to a venue is just the first step in realizing the potential revenue location-based Wi-Fi networks can provide. Organizations often lack insight into a user's behavior while in their venues. Understanding traffic patterns and dwell times is key to addressing user needs and improving operations and loyalty.

The CMX solution has three aspects:

- CMX Detect—Acknowledges a mobile consumer's presence in a venue by detecting the mobile device and its characteristics before they enter.
- CMX Connect—Provides premium mobile consumer access in a venue with seamless and secure Wi-Fi connectivity, allowing mobile consumers to receive personalized and location-based services. Organizations can collect these preferences and device and roaming credentials through direct access to a venue's network or through social media sites.
- CMX Engage—Organizations can gather highly relevant content and services based on user attributes and real time location to deliver a personalized, context-aware experience to a mobile consumer while in their venue.

The Cisco CMX solution relies on a Wi-Fi infrastructure within a venue. At the heart of the CMX solution are location-based services (LBS), which are essential to understanding a mobile user's context—where they are and what they are doing—and can help organizations engage with their users in a relevant way.

Cisco's CMX solution allows venues to simultaneously provide users with highly personalized content, provide services to customers to increase the customer experience, and gain visibility into customer behavior in their venues. CMX detects in-venue Wi-Fi enabled devices, prompts customers to connect to the wireless network, and engages them with value-added content and offers.

## CMX and the Wireless Infrastructure

The CMX solution relies on a Wi-Fi infrastructure within a venue—the key enabler for service delivery—to detect, connect, and engage with mobile users. An organization's WLAN network must become as robust, secure, scalable, and predictable as possible to ensure a positive experience for mobile users within a venue. Many venues do not have a wireless network since network access was not a key design factor during construction. Before the mobility explosion, venues used paper advertisements, maps, and brochures to provide users with information about their location. With the increase in users with smart devices, a Wi-Fi network is critical to provide services customized for the individual. This new way of engaging with users also saves organizations the money and resources required to publish and update paper products. When establishing a Wi-Fi network within a venue, critical design

considerations include determining how many access points are needed within a venue, guaranteeing that the wireless signal is sufficiently strong, and validating that there are no rogue APs that compromise the security of an organization and its venues.

## Role of the Mobile User in the CMX Network

The success of CMX requires establishing a mobile guest Wi-Fi network in the venue and a realization that this network is different from the organization's corporate network. Mobile users are often not employees of a venue and typically are only accessing data from the Internet or social media sites. Depending on the venue, there may be more mobile guest users than employees. This increase in smart devices used by mobile guests on a Wi-Fi network must be taken into consideration as increasing amounts of data traverse a venue's wireless network. In addition, the traffic from mobile guest users must remain separate from the organization's corporate network so the different types of traffic can be managed differently if network congestion occurs. Organizations must also protect their corporate data from mobile guest users to prevent them from injecting threats into the organization's network. The benefits of CMX to an organization include increasing the engagement and the loyalty of mobile users in their venues.

## Role of the Organization in the CMX Network

With the increase of mobile devices and Wi-Fi connectivity, CMX provides expanded opportunities for organizations to connect and engage with their users. Services and benefits can be sent to a mobile user's smart device depending on their location within a venue to ensure the experience within the venue is personalized and relevant to the user. Leveraging the Wi-Fi network and adding location based services to their network, organizations have the opportunity to use location analytics to understand how many users are in their venue, how long they dwell in a certain zone in their venue, as well as what paths they take while in the venue. This information is critical to an organization to provide the right engagement strategies at the right time in the right place. Location analytics are a valuable piece of the CMX solution that provides organizations the data they need to connect and engage with their mobile users while in their venue.

## Concerns for CMX and Mobile Users

- **Privacy**—In many organizations, in exchange for free Wi-Fi, a mobile user must accept the terms and conditions set by the venue. To receive Wi-Fi services, social media analytics or consumer information can be collected. While the advantages of a mobile user “opting-in” to the terms and conditions provides benefits such as services, discounts, and the right to use Wi-Fi, many mobile users prefer not to allow that aspect of their privacy to be invaded. This privacy concern is lessening as more and more mobile users check-in their location on social media sites, but it is still a valid concern. The workaround to this privacy concern is to not opt-in to the Wi-Fi network and use a cellular network for access to the Internet. The CMX solution provides a choice that accommodates all levels of privacy.
- **Security**—While there are privacy concerns regarding mobile user's rights, organizations are facing increasing amounts of guest traffic on their Wi-Fi networks. Although a wireless network may have been established for corporate users only, the addition of mobile users on the wireless network requires that organizations protect their data from mobile users and mitigate against threats to the organization. Having a strategy for protecting an organization's corporate data while providing valuable personalized data for their mobile users is critical. Strategies might include creating



mobile-user only wireless network access or offloading all mobile user data to the Internet. The CMX solution offers many design options to ensure an organization maintains the security of their critical assets while providing valuable services to their mobile users.

## CMX Solution Advantages

The CMX Solution provides benefits to both organizations and their mobile users, including:

- **Location-based services**—Organizations can use signals from Wi-Fi enabled devices to detect each user's location. They can also deliver location-based product information, offers, and ads to the user's device, with an option for a customized loyalty application.
- **Seamless Wi-Fi onboarding**—Recognizing the aforementioned requirement for a balance between privacy and security, the CMX solution supports easy opt-in onboarding of the user's mobile device while maintaining corporate security policies. With a simple touch of a screen, users can receive controlled access to the venue's Wi-Fi network.
- **Advanced analytics**—Organizations can gain insights into user's traffic patterns and trends through location analytics gathered from mobile device signals. Detailed reports can provide venues with valuable information on dwell times, traffic patterns, new versus repeat customers, and conversion rates on marketing campaigns. Data analytics can be used to drive operational efficiencies and improve customer service.
- **Targeted advertisements and messaging**—Organizations have a new opportunity to meet users' needs and preferences with personal and contextual offers that are based on the user's traffic patterns. Personalized pop-up messages, based on current location, can be delivered to the user's smart phone via a mobile app.
- **Indoor directions and venue services**—On a user's smartphone or tablet, organizations can display a list of primary departments or areas of interest on a virtual map as well as provide any services within these areas that might interest a mobile user.

The Cisco CMX solution is a proven solution design that is fully tested and documented in a Cisco Validated Design (CVD). The CVD program consists of solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. The Cisco CMX CVD integrates Cisco products, third-party products, and devices into a comprehensive approach to deploying CMX that provides these benefits to not only organizations, but to the mobile users they serve.



## **PART 2**

### **CMX Design Overview**





## Summary of CMX Design Overview

---

**September 4, 2014**

This part of the CVD describes the various components and technologies necessary for the implementation of a successful CMX solution. Multiple deployment models in which the components may be implemented within campus and branch network infrastructures are discussed. Finally security considerations as well as other additional considerations are discussed.

This part of the CVD includes the following chapters:

- **CMX Solution Components**—Highlights the wireless (Wi-Fi) network infrastructure necessary for providing location services and CMX services within this design guide. A discussion of the Cisco Context Aware Service (CAS), which provides location services, and the technologies behind CAS are also presented. Finally, an introduction to the various CMX services which make use of the location information provided by CAS is discussed.
- **CMX Deployment Models**—Introduces high-level models for the deployment of infrastructure components necessary for location services and CMX. Considerations around bandwidth utilization and scalability of the MSE are discussed. Finally the high-level models are mapped to campus and branch designs showing physical infrastructure designs for supporting CMX services as well as guest access for CMX Visitor Connect.
- **CMX Security Considerations**—Focuses on traffic isolation for guest wireless access as part of CMX Visitor Connect. Additionally, it discusses Role-Based Access Control (RBAC) for the Mobility Services Engine (MSE) as well as the CMX Connect & Engage service.
- **CMX Additional Considerations**—Highlights additional considerations when deploying a CMX solution, including:
  - How fast location information is updated and made available.
  - Considerations around specific mobile device platforms, such as Apple IOS 8 devices and some Android devices.
  - Considerations around the use of 2.4 and 5 GHZ frequency bands when deploying location services and CMX services.
  - Considerations around the deployment of the FastLocate feature.





# CHAPTER 3

## CMX Solution Components

---

September 4, 2014

This chapter highlights the wireless (Wi-Fi) network infrastructure necessary for providing location services and CMX services within this design guide. A discussion of the Cisco Context Aware Service (CAS), which provides location services, and the technologies behind CAS are also presented. Finally, an introduction to the various CMX services which make use of the location information provided by CAS is discussed.

## Wireless Infrastructure

The underlying infrastructure behind all CMX applications and services discussed within this design guide is the Cisco wireless LAN (IEEE 802.11) network infrastructure, which consists of the following hardware:

- Cisco Aironet Access Points (APs)
- Cisco Wireless LAN Controllers (WLCs)
- Cisco Mobility Services Engine (MSEs)
- Cisco Prime Infrastructure

## Cisco Aironet Access Points

Cisco Aironet access points provide Wi-Fi connectivity to the network infrastructure. Within this version of the CMX design guide, APs also assist in providing the following services:

- Location services for mobile devices—[Cisco Context Aware Service \(CAS\)](#) provides additional details around how APs participate in providing location services for mobile devices.
- Network connectivity for guest mobile devices.

The Cisco second generation APs in this design guide include the Cisco Aironet 3700, 2700, 3600, and 2600 Series.

Cisco 3700 Series APs are ideal for high-density network environments that use mission-critical, high-performance applications. They feature the industry's first AP with an integrated 802.11ac Wave 1 radio supporting a 4x4 multiple input, multiple output (MIMO) design with three spatial streams for data rates up to 1.3 Gbps. The flexible, modular design of the Cisco 3700 Series provides expansion capability for a future 802.11ac Wave 2 module and advanced services such as the Wireless Security Module (WSM).

Cisco 2700 Series APs are non-modular dual band (5 GHz and 2.4 GHz) 802.11ac access points optimized for adding capacity and coverage to dense Wi-Fi networks. They feature a 3x4 MIMO design with three spatial streams for a maximum data rate up to 1.3 Gbps.

The Cisco 3700 and 2700 Series APs incorporate the Cisco High-Density Experience (HDX), which includes among other features Cisco CleanAir® with enhanced support for 80-MHz channels and updated ClientLink 3.0 with support for 802.11a/b/g/n/ac. Cisco CleanAir® technology is enabled in hardware for both the Cisco 3700 and 2700 Series APs. Cisco ClientLink 3.0 helps improve performance of clients on the wireless LAN (WLAN).

Cisco 3600 Series APs are ideal for customers looking for best-in-class performance in 802.11n environments with high client density. They feature the industry's first 802.11n 4x4 MIMO design with three spatial streams for data rates up to 450 Mbps. The flexible, modular design of the Cisco 3600 Series provides expansion capability for emerging technologies such as the 802.11ac Wave 1 module and advanced services such as the WSM.

Cisco 2600 Series APs are dual band (5 GHz and 2.4 GHz) 802.11n access points ideal for mid-market small, mid-size, or large enterprise customers looking for mission critical performance. They feature a 3x4 MIMO design with three spatial streams for data rates up to 450 Mbps.

The Cisco 3600 and 2600 Series access points support additional technologies, such as Cisco ClientLink 2.0 and Cisco CleanAir®. Cisco CleanAir® technology is also enabled in hardware for both the Cisco 3600 and 2600 Series APs.

The field-upgradeable Wireless Security Module (WSM) has a dedicated dual-band radio with its own antennas enabling 7x24 scanning of all wireless channels in the 2.4 and 5 GHz bands. It offloads concurrent support for monitoring and security services—such as Cisco CleanAir® spectrum analysis, WIPS security scanning, rogue detection, context-aware location, and Radio Resource Management (RRM)—from the internal client/data serving radios within the Cisco 3700 or 3600 Series AP to the WSM. The WSM is required to enable the FastLocate feature (also known as All Packet RSSI or Data RSSI) for improved location currency. [Probe Request RSSI versus FastLocate](#) provides further details around the FastLocate feature.

**Note**

The Cisco 3700 Series AP requires 18 Watts and the Cisco 3600 Series AP requires 17 Watts of power with the WSM module. When powering the AP from a Cisco Catalyst switch, the switch port must support either POE+ (IEEE 802.3at standard) which supplies up to 30 Watts or Cisco Universal Power over Ethernet (UPoE) which delivers up to 60 Watts of power per switch port.

Cisco Aironet APs can operate as lightweight or autonomous access points. When functioning as lightweight APs, a wireless LAN controller (WLC) is required. In this design, the 802.11 MAC layer is essentially split between the AP and the WLC. The WLC provides centralized configuration, management, and control for the access points. All designs in this design guide assume lightweight APs.

Further information regarding Cisco Aironet APs can be found in the following at-a-glance document:

[http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps10981/at\\_a\\_glance\\_c45-636090.pdf](http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps10981/at_a_glance_c45-636090.pdf)

**Note**

Cisco Meraki wireless LAN infrastructure is not discussed within this version of the CMX design guide.



## Cisco Wireless LAN Controllers (WLCs)

Cisco wireless LAN controllers (WLCs) automate wireless configuration and management functions and provide visibility and control of the WLAN. Within this version of the CMX design guide, WLCs also assist in providing the following services:

- Location services for mobile devices—[Cisco Context Aware Service \(CAS\)](#) provides additional details around how WLCs participate in providing location services for mobile devices.
- Network connectivity for guest mobile devices.

Cisco WLC functionality can be within standalone appliances, integrated within Catalyst switch products, or run virtually on the Cisco Unified Computing System (UCS). The Cisco wireless LAN controller platforms included within this version of the CMX design guide include the Cisco 5508 WLC and the Cisco Flex 7510 WLC. The Cisco 5508 and Flex 7510 platforms run Cisco Unified Wireless Network (CUWN) software (also referred as AireOS software). The Cisco 5508 WLC is targeted for mid-sized and large single-site enterprises. Within this design guide it is deployed within the campus supporting APs operating in centralized (local) mode. The Cisco Flex 7510 WLC is targeted for enterprise branch environments. Within this design guide it is deployed as a remote controller supporting APs operating in FlexConnect mode. [Campus and Branch Designs](#) in [Chapter 4, “CMX Deployment Models”](#) provides details about the deployment of these WLC platforms.

[Table 3-1](#) shows scalability of these platforms in terms of APs, clients, and throughput.

**Table 3-1** *Wireless LAN Controller Scalability*

Platform	Access Points Supported	Clients Supported	Throughput
Cisco 5508	Up to 500	Up to 7,000	Up to 8 Gbps
Cisco 7510	Up to 6,000 APs with up to 2,000 FlexConnect groups	Up to 64,000	Up to 1 Gbps centrally switched traffic

Further information regarding Cisco WLC platforms can be found in the following at-a-glance document:

[http://www.cisco.com/en/US/prod/collateral/modules/ps2706/at\\_a\\_glance\\_c45-652653.pdf](http://www.cisco.com/en/US/prod/collateral/modules/ps2706/at_a_glance_c45-652653.pdf)

## Cisco Mobility Services Engine (MSEs)

The Cisco Mobility Services Engine (MSE) is a platform that helps organizations deliver innovative mobile services and improve business processes through increased visibility into the network, customized location-based mobile services, and strengthened wireless security. The following mobility services are supported on the MSE:

- Context Aware Service
- Wireless Intrusion Prevention System (wIPS)
- CMX Analytics (Location and Presence)
- CMX Connect & Engage (includes the Web service for guest access and the Cisco SDK for app development)
- Mobile Concierge Service

Mobility services are supported based upon the licensing of the MSE as shown below:

- Base Location Services (also called the Context Aware Service)—Requires Location Services licensing.
- Wireless Intrusion Prevention System (WIPS)—Requires WIPS licensing.
- CMX Analytics, CMX Connect & Engage, and the Mobile Concierge Service—Requires Advanced Location Services licensing.

This version of the Cisco CMX design guide discusses the following services:

- Location services for mobile devices—[Cisco Context Aware Service \(CAS\)](#) provides additional details around how the MSE participates in providing location services for mobile devices.
- CMX Analytics (both Location Analytics and Presence Analytics)
- CMX Visitor Connect (part of CMX Connect & Engage)

The Cisco MSE is available as a physical appliance or as a virtual appliance. Additional information regarding the Cisco MSE platform can be found in [MSE Scalability](#) in [Chapter 4, “CMX Deployment Models.”](#)



**Note**

Use of the Cisco SDK for mobile app development will be discussed in future versions of the Cisco CMX design guide.

## Cisco Prime Infrastructure

Cisco Prime Infrastructure (PI) is the continued evolution of Cisco Prime Network Control System (NCS). It interacts with Cisco wired and wireless infrastructure components to be a central management and monitoring portal. Cisco PI configures and monitors Catalyst switches and routers and it also controls, configures, and monitors all wireless LAN controllers (WLCs) and, by extension, all access points (APs) on the network.

Within the Cisco CMX design guide, Cisco Prime Infrastructure also provides the following services:

- Provides the administrative interface for importing and tuning floor maps for location services.
- Integrates with the Cisco MSE to synchronize floor maps.
- Synchronizes MSE services with WLCs.
- Integrates with CMX Presence Analytics to import APs not associated with any floor map.
- Provides the administrative interface for enabling MSE services such as the Context Aware Service (CAS), CMX Analytics, and CMX Visitor Connect.


Multiple WLCs and MSEs may be managed and monitored by Cisco Prime Infrastructure. Detailed information regarding floor maps and enabling MSE services via Cisco PI is provided in [Chapter 24, “Configuring Cisco Prime Infrastructure.”](#)

## Cisco Context Aware Service (CAS)

The underlying technology behind CMX applications and services is wireless (Wi-Fi) location. Location services are provided to Cisco wireless network infrastructures through the Context Aware Service (CAS) running on the Cisco MSE. CAS provides the location database which is leveraged by CMX applications and services.

Figure 3-1 provides a high-level overview of the information flows between the various hardware components for CAS. It assumes the WLAN within the site has been designed for location services. Detailed information regarding designing the WLAN within the site to support location services is provided in the “CMX Radio Frequency and Location Based Design” part of this design guide.

Figure 3-1 Context Aware Service (CAS) Hardware and Data Flows



Each of the steps in Figure 3-1 is explained below:

- Step 1** To be able to collect Wi-Fi location information for a site, the network administrator must first set up the wireless infrastructure to support location services, which includes:
- Importing the floor map for the site into Cisco Prime Infrastructure.
  - Correctly sizing and tuning the floor map.
  - Placing APs in the correct location on the floor map.
  - Enabling the Context Aware Service (CAS) on the MSE.
  - Syncing the WLC and MSE through Cisco Prime Infrastructure.

The network administrator accomplishes this by establishing an HTTPS session to the Cisco PI server and using the graphical user interface (GUI).

The network administrator must also synchronize the floor map information with the MSE, which pushes the floor map information to the MSE. The interface between Cisco PI and the MSE uses SOAP/XML & REST messages over HTTPS.

The network administrator must also synchronize MSE services like CAS with the WLC so that the WLC forwards collected data (location, intrusion detection, etc.) from the APs to the MSE.

Additional configuration of the Context Aware Service (CAS), CMX Presence Analytics, CMX Location Analytics, and CMX Visitor Connect must also be done by directly establishing an HTTPS session to the MSE running the associated service.

Detailed information regarding setting up the wireless infrastructure to support location services is provided in the “[CMX Configuring the Infrastructure](#)” part of this design guide.

- Step 2** FastLocate Only—If the deployment is using the FastLocate feature, the APs need to be time synchronized via NTP so that Wireless Security Modules (WSMs) all simultaneously scan the same channel as they proceed through the scan list. Detailed information regarding the differences between the Context Aware Service using Probe Request Received Signal Strength Indication (RSSI) and the FastLocate feature is provided in [Probe Request RSSI versus FastLocate](#).
- Step 3** For the Context Aware Service (CAS) to function, wireless clients must either send Probe Requests on each active channel or associate with an AP and send packets if using the FastLocate feature.
- Step 4** Each AP within range of the wireless client either hears Probe Requests sent by the wireless client on the 2.4 GHz channel and/or 5 GHz channel on which the AP is operating or hears packets when the monitoring radio within the WSM module dwells on the channel on which the wireless client is operating on—when the wireless client is associated to an AP and the FastLocate feature is enabled. RSSI information is calculated for the particular client from either the Probe Requests or from packets sent by the wireless client.

A minimum of three APs are needed to determine the X,Y coordinates of the wireless client relative to the floor map. However accuracy is highest when a wireless client is seen by at least four APs. For wireless (Wi-Fi) locations, the APs must be configured onto a floor map within Cisco PI, which is then synchronized with the MSE. If the AP has not been placed on the floor map which is synchronized with the MSE, RSSI information is still calculated by the AP and forwarded to the MSE. However the MSE does not use the RSSI information from the particular AP in determining the X,Y coordinates of the wireless client. If the RSSI values calculated by all of the APs and sent to the MSE are below the RSSI cutoff threshold setting within the MSE, the MSE ignores the calculation and the data point is not stored in the MSE location database. By default this is set for -65 dBm. Information about setting this parameter is provided in [Chapter 25, “Configuring the Mobility Services Engine for CMX.”](#)

For wireless (Wi-Fi) presence, APs which do not appear on floor maps must be imported to the MSE and associated with a Presence site. Information showing how to do this is provided in [Chapter 26, “Configuring CMX Analytics.”](#)

If only one AP sees the wireless client and the RSSI value is above the RSSI cutoff threshold, the location of the wireless client is reported to be the X,Y coordinate of the AP itself, relative to the floor map.

Each AP aggregates messages which contain RSSI information and sends them to the WLC which controls the AP approximately every 500 milliseconds via the CAPWAP protocol.

- Step 5** The WLC aggregates RSSI information for each client from each AP which it controls and forwards all messages to all MSEs—every two seconds by default—using the Network Mobility Services Protocol (NMSP).
- Step 6** Since RSSI information regarding a given wireless client could come from APs on the same floor map, but controlled by different WLCs, the location services (CAS) engine within the MSE aggregates data for five seconds before calculating locations of wireless clients. Once the location calculation is completed, the MSE can update the location (CAS) database for the particular wireless client. Services such as the CMX Analytics engine within the MSE can then make use of the updated information within the location database of the MSE.

**Note**

CMX Analytics has its own database, which is built off of the MSE location database. CMX Analytics periodically pulls information (in batch mode) from the MSE location database. Therefore CMX Location Analytics should not be used for real-time analysis.

## Probe Request RSSI versus FastLocate

Prior to WLC and MSE release 8.0, the location services engine within the Context Aware Service (CAS) relied solely on IEEE 802.11 Probe Requests to calculate the location of wireless clients via RSSI information. Probe Requests are sent when the wireless client actively scans for a Basic Service Set (BSS)—in other words an Access Point (AP) with which to join. Probe Requests are good candidates for collecting RSSI information because the wireless client typically probes multiple channels to develop a scan report which is then used by the client to select which BSS/AP to join. Wireless clients typically cycle through 5 GHz and 2.4 GHz channels as they send Probe Requests, waiting for Probe Responses.

**Note**

Probe Requests are not sent simultaneously to all active channels. The wireless client typically sends a probe request on a particular channel and briefly listens for a Probe Response before switching channels and sending another probe request.

In [Figure 3-1](#), three APs operating in both the 2.4 GHz and 5 GHz frequency bands are shown. As the wireless client generates Probe Requests on channels 36, 40, and 44 in the 5 GHz frequency band, each of the APs hears, respectively, the Probe Request.

Unfortunately client probing frequency in most smartphone and tablet devices has been decreasing over time and is also non-deterministic. Probe Request frequency can vary from under a second to five minutes depending on the smartphone or tablet device operating system, wireless driver, current activity on the device, battery usage, etc. Active scanning consumes battery power of mobile devices. Some smartphones disable active scanning altogether below a certain percentage of remaining battery power. Hence, such devices are virtually non-trackable when remaining battery power drops below a certain threshold. Additional information is discussed in [Chapter 6, “CMX Additional Considerations.”](#)

Because of the non-deterministic nature of Probe Requests, total location error—which is a function of location accuracy and location currency—is increased. A detailed discussion of the expected location accuracy and location currency of the Cisco Context Aware Service (CAS) is provided in [Chapter 13, “Location Fundamentals.”](#)

To alleviate the issue of decreased location currency, Cisco introduced a new feature called FastLocate. FastLocate is implemented in WLC and MSE version 8.0 and requires Prime Infrastructure release 2.1 to enable it. Information regarding how to enable the FastLocate feature is provided in [Chapter 23, “Configuring Cisco Wireless LAN Controllers.”](#)

With the FastLocate feature enabled, RSSI information is collected on all data packets transmitted by the wireless client, not just Probe Requests. The FastLocate feature is intended to make the collection of RSSI data more deterministic and more current to reduce movement error, therefore resulting in a reduction in the total location error. The FastLocate feature requires Cisco 3600 or 3700 Series APs with the Wireless Security Module (WSM). The WSM provides a separate, dedicated dual-band radio which allows the AP to monitor other channels for CleanAir or wIPS purposes and simultaneously service data from wireless clients. With the FastLocate feature, WSM channel scanning is synchronized across APs using Network Time Protocol (NTP). The result is that all WSMs within the site listen to the same 5 GHz or 2.4 GHz channel at the same time.

The frequency by which each channel is monitored (also referred to as the return time to channel) influences the currency of location information (also referred to as the location refresh rate) obtained by FastLocate. The frequency by which each channel is monitored is based upon the number of channels within the FastLocate scan list and the dwell time (Tdwell) on each channel. The scan list (also referred to as the off-channel scan list) is the list of channels to be monitored for activity. The dwell time is the amount of time the WSM radio monitors that particular channel along with the time required to change channels. Whether CleanAir is enabled or disabled influences the return time to channel because instead of the channel list being just the channel slots for FastLocate, it has the channels slots for CleanAir as well. Hence the return time to channel is longer when CleanAir is enabled.

The following provides an example based on a U.S. deployment with both 2.4 and 5 GHz (U-NII-1, U-NII-2 non-extended, and U-NII-3 channels) operation with CleanAir enabled.

Channels: 2.4GHz non-overlapping U.S. country channels: 1 6 11

5 GHz U.S. country channels (16 channels, excluding U-NII-2 extended): 36 40 44 48 52 56 60 64 149 153 157 161 165

Dwell time per FastLocate Slot = 250 milliseconds

Dwell time per CleanAir Slot = 175 milliseconds

Based on the above parameters, the WCM scan list would be as follows:

1, 6, 11, 36, X, 40, 44, 48, 52, X, 56, 60, 64, 149, X, 153, 157, 161, 165, X

Each of the numbers (1, 6, 11, ...) indicate 250 milliseconds of time that FastLocate dwells on that particular channel collecting RSSI information from wireless devices operating on that channel. Every fifth slot shows an X. Each X indicates 175 milliseconds of time that CleanAir dwells on additional channels not included in the FastLocate scan list.

Given this example, the return time to channel for FastLocate would be estimated as follows:

16 FastLocate channels \* 250 milliseconds of FastLocate dwell time per channel = 4,000 milliseconds

Plus

4 CleanAir channels \* 175 milliseconds of CleanAir dwell time per channel = 700 milliseconds

4,000 milliseconds + 700 milliseconds = 4,700 milliseconds or approximately 5 seconds.

This allows for both the FastLocate feature and CleanAir to operate simultaneously on the Wireless Security Module (WSM). Increasing the number of channels in the scan list increases the return time to channel. Decreasing the number of channels in the scan list decreases the return time to channel.

The FastLocate feature requires the wireless device to be associated to and communicating with an AP to take advantage of increased currency of location information. To be seen during every scan cycle, which is the optimal currency of location information using the FastLocate feature, the wireless client must be transmitting packets during the time the WCM modules dwell upon the channel which the wireless client is operating.



#### Note

CMX Location Analytics does not necessarily require association of the wireless device to any network. Hence total location accuracy of a CMX Location Analytics deployment which does not involve association of the wireless device to the network may not improve with FastLocate. However for CMX services which involve connecting the wireless device to a network, total location accuracy may improve due to increased currency of location information and additional data points upon which to base location calculations.

One way of accomplishing this is via an app running on the mobile device which transmits packets during every scan cycle, however this may not always be feasible to develop and deploy. Hence FastLocate incorporates an additional feature which keeps track of unresponsive (idle) clients.

Unresponsive clients are devices associated to the wireless infrastructure, but which RSSI information has not been refreshed for a given number of scan cycles. By default this is 10 scan cycles or roughly from 40 to 60 seconds depending on the channel scan list and whether CleanAir is enabled. This is a configurable parameter. Information on configuring this parameter is provided in [Chapter 23, “Configuring Cisco Wireless LAN Controllers.”](#)

When a wireless client has been determined to be unresponsive, an 802.11 Block Acknowledgement Request (BAR) is sent to the wireless client by the AP to which the client is associated, shortly before scanning the channel to which the particular wireless client is associated, during the next scan cycle. The wireless client should respond to the BAR with a Block Acknowledgement (BA). This ensures that the particular wireless client is heard during that scan cycle and that the RSSI information for that particular wireless client is refreshed.

## Connected Mobile Experiences Services

The overall Cisco CMX solution can be broadly separated into three levels of functionality as discussed in [Chapter 1, “Connected Mobile Experiences Solution Overview”](#):

- **CMX Detect**—Detects the presence and/or location of a mobile device within a venue. This mobile device could belong to a customer within a retail establishment, a patient within a healthcare facility, a patron to a museum, a traveler within an airport, an employee within a corporate location, etc. Insight into movement, dwell times, and crowding within the venue can then be acquired to provide improved service.
- **CMX Connect**—Provides an easy-to-use and scalable method of connecting a mobile device to the guest wireless LAN network within a venue. This can be used to provide some level of context-based services to visitors while at the venue or to potentially gaining insight into the demographics of visitors to the venue via social media sites.
- **CMX Engage**—Provides context-based services to the visitor through their mobile device as they enter and move through various points-of-interest (POIs) in the venue.

These three levels of functionality are delivered via the following CMX services:

- **CMX Analytics** (includes Location Analytics and Presence Analytics)
- **CMX Connect & Engage** (includes CMX Visitor Connect, CMX Facebook Wi-Fi, the CMX Mobile Application Server, and the CMX SDK for mobile app development)

This version of the Cisco CMX design guide discusses CMX Location Analytics, CMX Presence Analytics, and CMX Visitor Connect. The following sections provide a high-level overview of each of the CMX services discussed within this version of the CMX design guide.

## CMX Location Analytics

Cisco CMX Location Analytics makes use of the location information collected by the Context Aware Service (CAS) running on the MSE to determine mobile device parameters such as:

- **Dwell time**—How long people stay in a specific point.
- **Crowding**—Popular points at which people stay a long time.
- **Path choice**—For example, do people usually turn left or right when coming out of an elevator.

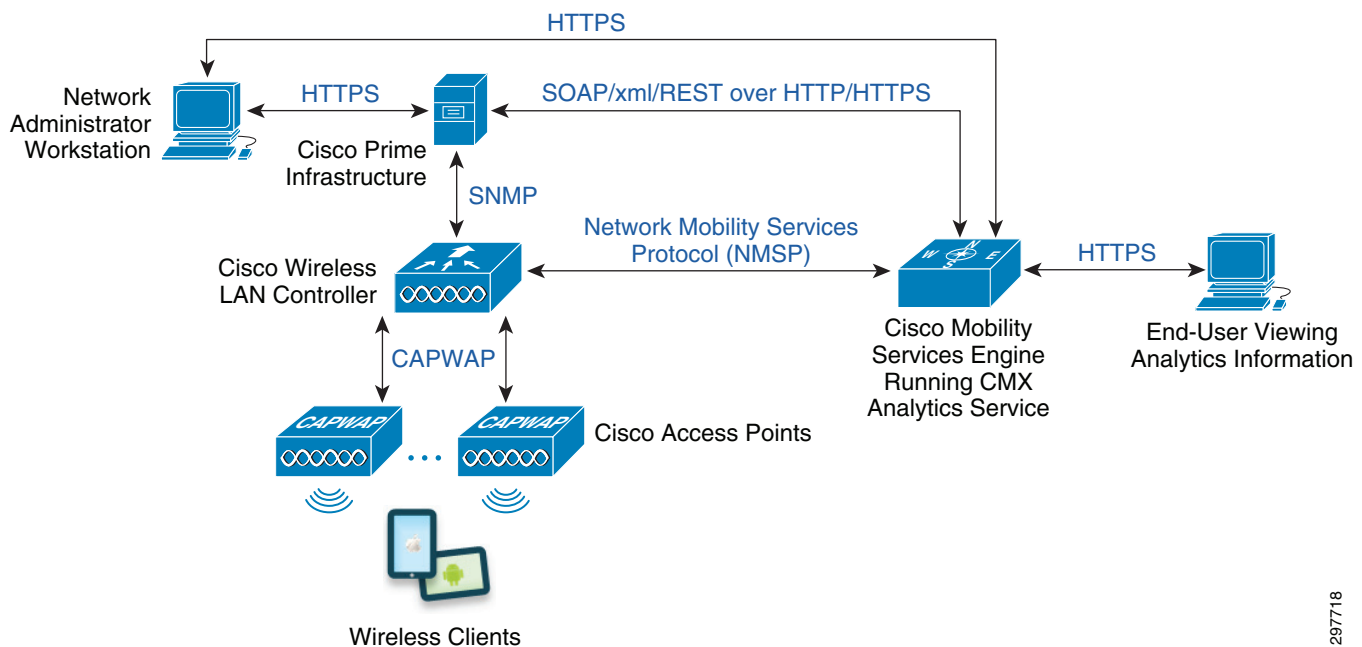


CMX Location Analytics aggregates this information for common understanding, so businesses can use this information to better understand how their customers interact with different parts of their venues or environments. Businesses can utilize CMX Location Analytics to help achieve better facility planning, measure changes in their buildings, and improve their interaction with customers.

The basic data used by CMX Location Analytics is in the form of MAC addresses, time, X and Y coordinates, etc. CMX Location Analytics helps aggregate and visualize this data, consisting of anonymous MAC addresses, to help generate insights about the movement and behavior patterns of the people using mobile devices who are visiting a venue. This can be used to help provide better service to visitors of the venue. A venue can be a shop, mall, airport, or city center, provided that it has a network of wireless access points so that devices moving within that space can be located. All wireless (Wi-Fi) devices have unique MAC addresses which are used for communicating with the network infrastructure (access point). Without the use of MAC addresses, the wireless (Wi-Fi) network itself would not operate. All forms of networking utilize some form of unique addressing to deliver information to the correct device. Although there may be privacy concerns around the use of MAC addresses for analytics, it should be noted that MAC addresses are associated with a physical device and are not associated with any specific end-user information.

Figure 3-2 shows a high-level overview of the hardware and information flows for CMX Location Analytics.

**Figure 3-2 High-level Overview of Hardware and Information Flows for CMX Location Analytics**



297718

As can be seen, the hardware and information flows are basically the same as those shown in Figure 3-1 for the Context Aware Service (CAS). CMX Location Analytics is a separate process which can run on the same MSE that runs the Context Aware Service (location services) or on a separate MSE. [MSE Scalability](#) in Chapter 4, “CMX Deployment Models” discusses this further.



**Note**

CMX Location Analytics periodically (approximately every 15 minutes) extracts information from the Context Aware Service (location services) database running on the MSE to aggregate location information from various mobile devices for analysis and/or reporting. Hence the information presented within CMX Location Analytics is not real-time, but historical information. Therefore slight variations



can happen for the same analysis or report taken at different dates due to the historical data. On startup of the CMX Location Analytics service, it can be up to 45 minutes before analytics information is displayed.

CMX Location Analytics data is accessed by establishing an HTTPS session directly to the MSE which runs the Location Analytics service. The functionality of CMX Location Analytics is separated into three tabs located at the top of the CMX Analytics home page:

- Dashboard Tab
- Analytics Tab
- Reports Tab

Each is discussed below.

## Dashboard Tab

The Dashboard tab provides a quick and easy way of visualizing device counts and dwell times of devices in various zones and timeframes throughout the venue. The CMX Analytics Dashboard can be customized by the CMX administrator by adding or deleting pages and widgets, which include:

- For Location Analytics sites, whether device count or dwell time is displayed
- For Presence Analytics sites, whether device count, dwell time, or conversion percentage is displayed
- Whether the information is displayed in bar chart or line chart format
- The particular zone from the floor map to be included within the widget
- The start and end dates to be displayed
- The start and end times of the day to be displayed

Further details regarding the CMX Analytics Dashboard are provided in [Chapter 26, “Configuring CMX Analytics.”](#)

## CMX Analytics Tab

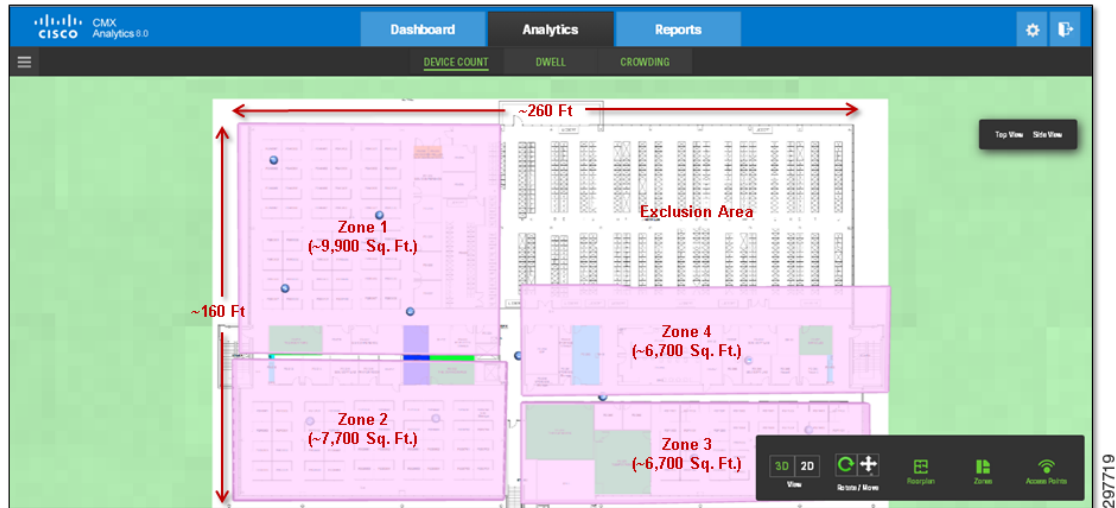
The CMX Analytics tab is used to perform deeper analysis on devices that pass through the venue. The results can be viewed in 3D or 2D within the environment. Various types of analysis can be performed, including:

- Zone Analysis—Provides aggregated parameters such as dwell time, number of devices, and crowding for each zone defined.
- Alternative Path Analysis—Shows a breakdown of the percentage of devices going to each destination from each starting point and vice versa.
- Heat Maps—Provides a graphical representation of point data which can be viewed on the map in such a way that areas of higher concentration appear darker.
- Typical Locations—Provides parameters such as dwell time, number of devices, and crowding for different areas of the building.

Zones within CMX Analytics represent separate areas within a venue where analytics information is aggregated. Zones are configured as coverage areas which are defined on a floor plan within Cisco Prime Infrastructure (PI). It is recommended that zones should be configured to be no smaller than

approximately 1,000 square feet. [Figure 3-3](#) provides an example of zones configured on a floor plan (as viewed through the CMX Analytics tab) within a venue. Note that the building dimensions and zone sizes have been added to the figure.

**Figure 3-3** Example of Location Analytics Zones



**Note**

When changing zone names or adding new zones, it should be noted that analytics data will only be available for that zone on and after the date on which the zone name was changed or the new zone was added. Hence the CMX administrator should carefully plan zones such that changing zone names and/or adding and deleting zones is kept to a minimum.

Further details regarding the configuration of zones are provided in [Chapter 24, “Configuring Cisco Prime Infrastructure.”](#)

## CMX Reports

The CMX Location Analytics reporting facility provides a more regular and manager-oriented set of information through the provision of parameterized templates to measure various common trends and patterns that occur over a period of time in a particular zone. The time window for CMX Location Analytics reports is typically longer than for CMX Location Analytics analysis, discussed in [CMX Analytics Tab](#).

The CMX Location Analytics Reports tab has the following reports:

- Conversion Percentage Report—Estimates the percentage of people who were in the vicinity of the actual zone before entering that zone.
- Detected vs. Connected Devices Report—Shows an overview of the number of devices that were connected to the network and the devices that were merely probing during a given time period for a particular zone.
- Daily Visitors and Dwell Times Report—Shows both the number of devices and the average time spent in a zone across several days.
- Hourly Visitors and Dwell Times Report—Shows both the number of devices and the average time spent in a zone across several hours.

- Movement between Zones Report—Provides a breakdown of all zones at specific points as devices pass to and from the focus zone.
- Repeat Visitors Report—Of the visitors who appeared within the venue within a defined time window (particular day, week, month, etc.), shows how frequently those same visitors returned to the venue since a defined start date (which can be before the time window).

Reports can be viewed directly on the MSE—requiring direct web access (HTTPS) to the MSE—or exported as an Adobe Acrobat (.pdf) file and manually emailed or printed.

## Differences in CMX Reports, Dashboard, and Analysis

CMX Reports and the CMX Dashboard operate off an aggregated (i.e., summarized) database which is smaller than the full CMX Analytics (analysis) database, which is restricted by default to 8 million points. This is to provide a fast response for the CMX Dashboard and CMX Reports. For this reason less history may be available for Analysis than for Reporting. Information within the aggregated database is updated approximately every hour.

CMX Reports and the CMX Dashboard also use a slightly different interpretation of a visit than CMX Analytics (analysis). For CMX Reports and the CMX Dashboard, a visit refers to a device seen that day at the site. If a device arrives at the zone or site, leaves for a few hours, and arrives at the zone or site again, it is still viewed as the device having visited the zone or site that day one time. For CMX Analytics (analysis), a visit to a zone refers to a device seen in that zone. If the device moves to another zone and is seen, then moves back to the original zone, CMX Analytics counts that as two visits to the original zone and 1 visit to the other zone. Alternatively, if the device is seen in a zone or site, is then not seen for more than hour, and re-appears in the zone or site, CMX Analytics counts that as two visits to the zone or site.

Because of this difference in the interpretation of visits between CMX Reports and CMX Dashboard and CMX Analytics (analysis), the information within each output may appear different even though the same zones or sites and dates were selected.

## CMX Presence Analytics

CMX Presence Analytics is targeted for customer locations with a small number (perhaps only one or two) of APs. These are referred to as “sites” from a Presence Analytics perspective. Hence sites are just logical groups of APs with the following attributes:

- A globally unique ID
- A globally unique name
- A description

In deployments with only one or two APs per site, CMX Location Analytics is of limited value since the infrastructure cannot track user movement with any degree of accuracy with only one or two APs. Further, there may be no need to track user movement in a small venue. However customers may still wish to collect analytics information, such as the percentage of passing customers who actually visit the site, the dwell times of customers within the site, and the crowding of the site during various times of the day. This information can be useful for staffing and promotional activities within the small venue.

The behavior of CMX Presence Analytics is as follows:

- If a wireless device is detected at a power level below the Low RSSI Threshold (default of -95 dB), the wireless client is discarded (ignored) from Presence Analytics.

- If a wireless device is detected at a power level above the Low RSSI Threshold, the wireless client is classified as a passer-by.
- If a wireless device is detected at a power level above the High RSSI Threshold (default of -75 dB), for a time period greater than the Dwell Time used to classify the wireless device as a visitor (default of 5 minutes), during the Time Period used to classify the wireless device as a visitor (default of 15 minutes)—then the wireless device is classified as a “visitor”.
- If a wireless device is detected at a power level above the High RSSI Threshold, for a time period less than the Dwell Time used to classify the wireless device as a visitor, during the Time Period used to classify the wireless device as a visitor—then the session is maintained. Note that the wireless device has already been classified as a “passer-by”.
- If a wireless device associates to the AP at the site, the wireless device is classified as a “visitor”.

Internally, Presence Analytics data is persisted to the internal MSE database. As with Location Analytics, Presence Analytics data should only be viewed as “off-line” analytics, meaning that the use cases should be historical only, not real-time. Presence statistics information includes:

- Site
- Number of visitors to the site
- Total devices seen
- Number of new visitors to the site
- Number of visits to the site
- Average dwell time at the site
- Average number of repeat visitors as of the past month

CMX Presence Analytics represents somewhat of a sub-set of the functionality of CMX Location Analytics, in that the analytics information, i.e. dwell times, device counts, and crowding, are based on the entire site versus individual floors or zones within a site.

## CMX Visitor Connect

CMX Visitor Connect is part of the CMX Connect & Engage service which runs on the MSE. CMX Visitor Connect is an easy-to-use method of connecting a mobile device to the guest WLAN within a venue utilizing the MSE and optionally social media sites. Along with guest Internet connectivity, CMX Visitor Connect provides the following functionality:

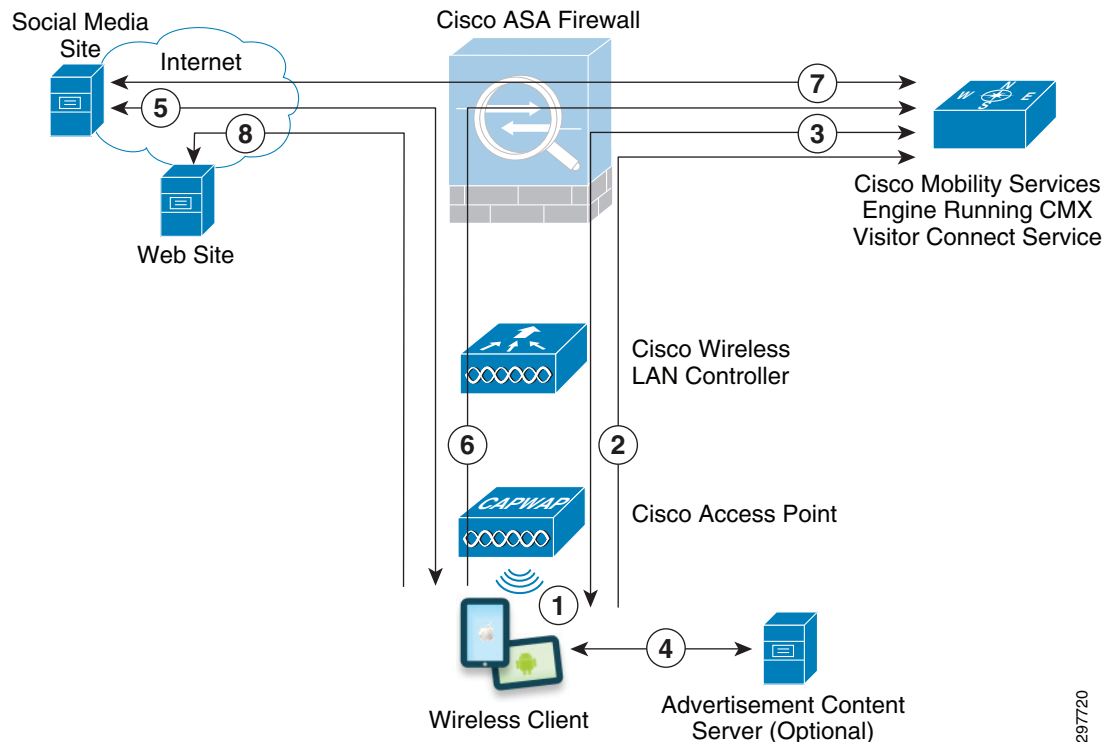
- The ability to push a customizable splash page to the guest mobile device during the login process, requesting basic information such as name and email address.
- The ability to push Terms & Conditions to the guest mobile device for use of the guest Wi-Fi service.
- The ability to authenticate the guest via any of the social media connectors currently supported by CMX Visitor Connect including Facebook, LinkedIn, and Google+.
- The ability to allow the guest to “opt-out” and still access the Internet without logging in to social media.
- The optional ability to push promotional/advertisement content to the guest device during the login process.
- The ability for the guest device to temporarily dis-associate to the guest Wi-Fi, then re-associate, without having to log in again.

- The optional ability to determine if a guest has logged in via social media or logged in anonymously and enforce different usage quotas based on that. This includes preventing a returning guest device from accessing the guest Wi-Fi network if it exceeds the daily usage quota and also isolating the device from the guest Wi-Fi network if it exceeds the daily usage quota.

Note that the splash page, Terms & Conditions, and promotional content are automatically sized appropriately for the mobile device type.

Figure 3-4 shows a high-level overview of the hardware and information flows for a guest mobile device utilizing CMX Visitor Connect to access the Internet.

**Figure 3-4 High-level Overview of Hardware and Information Flows for CMX Visitor Connect**



Each of the steps in the figure above is explained below:

- Step 1** The mobile device must first associate to an AP which broadcasts the B2C guest SSID / WLAN within the venue.



**Note** CMX Visitor Connect is often implemented to provide guest Wi-Fi and Internet connectivity for consumers visiting a venue (discussed in [CMX Visitor Connect Use Case Story](#) in [Chapter 7, “CMX Use Case Stories”](#)). Hence the guest WLAN is also referred to as the Business-to-Consumer (B2C) guest WLAN within this design guide to differentiate it from other types of guest Wi-Fi connectivity which may be implemented by an organization within a venue, such as corporate sponsored guest access. This is because different types of guest Wi-Fi connectivity may have different requirements for authentication, access control, etc.

The B2C guest WLAN is configured with no Layer 2 security (i.e., an open SSID) with Layer 3 Web Passthrough. Detailed information regarding configuration of the B2C Guest WLAN on Cisco WLCs for CMX Visitor Connect is provided in [Chapter 23, “Configuring Cisco Wireless LAN Controllers.”](#)

- Step 2** The end-user of the mobile device opens their web browser to reach a web site on the Internet. The web session is then redirected by the Cisco WLC to the CMX Visitor Connect service running on the MSE. The specific URL to which the guest mobile device is redirected is:

`http://<MSE_IP_Address>:8083/visitor/social.do`

MSE\_IP\_Address corresponds to the IP address of the MSE server. The CMX Visitor Connect service runs on a separate TCP port (8083) from other MSE services, such as the administrative web interface. Hence from a security perspective, it is recommended to limit mobile devices which associate to the B2C Guest WLAN to only reach TCP port 8083 of the MSE. The designs shown within this design guide assume that B2C guests are terminated outside of a Cisco ASA firewall, which provides stateful access control for B2C guest mobile devices. [B2C Guest Access for CMX Visitor Connect](#) in [Chapter 4, “CMX Deployment Models”](#) discusses the network infrastructure design and ASA firewall policy in more detail.

- Step 3** On re-direction of the web session, CMX Visitor Connect presents the end-user with a splash page for registration, terms & conditions, and the option for the end-user to login via social media sites.

The splash page is customizable and can be used to collect information such as the Name and Email Address of the visitor. Multiple splash pages can be configured within Visitor Connect, each used for a different venue in which the visitor is logging in or even for different points of interest (POIs) within a single venue. This is because the Context Aware Service, also running on the MSE, is aware of the location of the mobile device based on one or more of the following:

- When the mobile device generates Probe Requests, for instance before associating with the AP.
- When the mobile device associated with the AP.
- When the mobile device generates packets and if the FastLocate feature is enabled for the particular venue.

- Step 4** Optional: Optionally, CMX Visitor Connect can be configured to present marketing content, such as promotional ads, coupons, etc. to the mobile device. The actual content may be located on another server which must be accessible to the mobile device.

- Step 5** If CMX Visitor Connect is configured to allow the end-user to log in via one of the social media connectors supported (Facebook, LinkedIn, or Google+) and if the end-user chooses to log-in via social media site, the web session is redirected to the social media site. The end-user then enters their credentials for the particular social media site. This requires the CMX administrator to have previously configured the CMX Visitor Connect connector for the particular social media site. CMX Visitor Connect uses a variation of the OAuth 2.0 protocol for authentication to the Wi-Fi network via social media site.

- Step 6** On entering user credentials in the social media site, the browser of the mobile device is again re-directed back to the CMX Visitor Connect service running within the MSE, along with an authorization code.

- Step 7** The CMX Visitor Connect service running on the MSE authenticates that the end-user logged into the social media site by sending the authorization code, along with a ClientID and a Secret which was previously obtained during the configuration of the social media connector within the CMX Visitor Connect service on the MSE.

The CMX Visitor Connect service running on MSE software release 8.0 supports two user groups—the Social user group and the Basic user group. These groups are used to distinguish whether a mobile device’s end-user has logged in via a social media site or has accessed the guest WLAN anonymously. Note that the anonymous access to the guest WLAN requires the CMX administrator to enable that option within the CMX Visitor Connect service on the MSE.

The distinction between user groups can be used to provide different quota limits for the amount of traffic which a mobile device can send or receive per day based upon whether the device is placed into the Social user group or the Basic User group. If a quota is configured for a particular user group and if the mobile device exceeds that quota for the day, the mobile device is disassociated from the AP. If the mobile device reconnects to the guest WLAN, any web sessions are automatically redirected to a web page within the MSE indicating that the end-user has exceeded their quota for the day. Note that the time interval (daily from midnight to midnight) is not currently configurable and is based on the time zone of the MSE itself, not the time zone of the particular venue.

The MSE enforces quota limits by periodically collecting traffic information from the WLCs for wireless devices which have quota limits set (identified by the MAC address of the device) through the NMSP protocol. Removal of a device from the WLAN is initiated by the MSE authorizing disassociation of the wireless client, also through the NMSP protocol.

- Step 8** After authenticating that the end-user logged into the social media site, the CMX Visitor Connect service running within the MSE re-directs the web browser of the mobile device to the original site which the end-user was attempting to reach. Alternatively, CMX Visitor Connect can redirect the web browser of the mobile device to a site pre-determined by the CMX administrator.
- 

Detailed information regarding configuration of the MSE for CMX Visitor Connect is provided in [Chapter 27, “Configuring CMX Visitor Connect.”](#)

CMX Visitor Connect is dependent upon the wireless infrastructure design for guest Internet access. [B2C Guest Access for CMX Visitor Connect](#) in [Chapter 4, “CMX Deployment Models”](#) presents a method of providing guest wireless access using an anchored wireless LAN controller design and discusses an example of the configuration of the Internet Edge ASA security appliance interfaces for supporting CMX Visitor Connect.







# CMX Deployment Models

---

September 4, 2014

This chapter introduces high-level models for the deployment of infrastructure components necessary for location services and CMX. Considerations around bandwidth utilization and scalability of the MSE are discussed. Finally the high-level models are mapped to campus and branch designs showing physical infrastructure designs for supporting CMX services as well as guest access for CMX Visitor Connect.

## Overview

This section of the Cisco CMX CVD is targeted for IT personnel who are looking at deploying CMX services over private Enterprise network infrastructures. The common characteristic of Enterprise networks are that they are deployed for private use by a single business entity.

This version of the CMX CVD focuses primarily on meeting the CMX Location Analytics use cases, CMX Presence Analytics use case, and the CMX Visitor Connect use case, which is discussed in [Chapter 7, “CMX Use Case Stories.”](#) Each of these use cases is deployed over an enterprise wireless network infrastructure which is designed to support location services.



**Note**

---

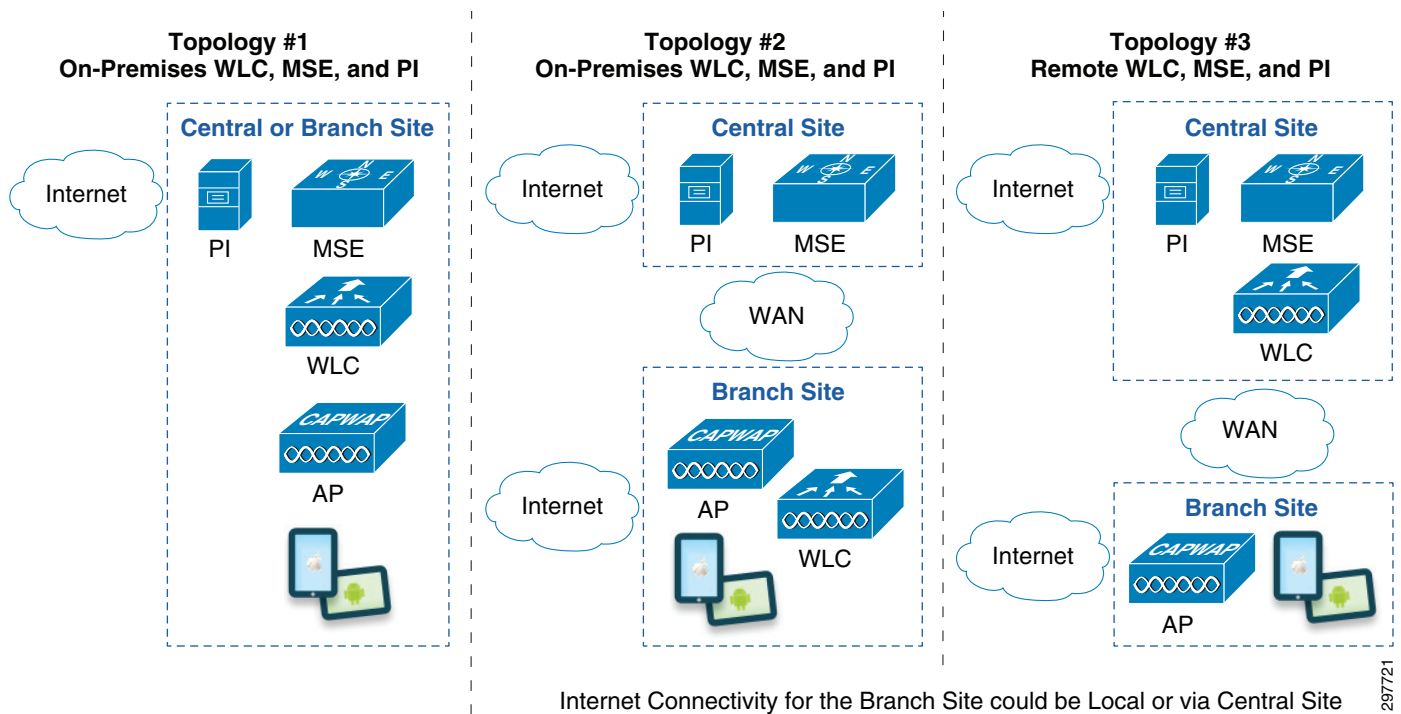
This version of the CMX design guide does not discuss deployment of CMX or location services over a Service Provider network infrastructure.

---

## Deployment Topologies

[Figure 4-1](#) shows three basic topologies for deployment of the components discussed in [Chapter 3, “CMX Solution Components”](#) in a network infrastructure.

Figure 4-1 High-Level Deployment Topologies



These topologies are based on the physical position of the Mobility Services Engines (MSEs) and Cisco Prime Infrastructure (PI) in relation to the wireless LAN controllers (WLCs) and access points (APs). The location of the APs represents the site where CAS (location services) and CMX services are required.

In Topology #1, all wireless infrastructure components—APs, WLCs, MSEs, and PI—are deployed on-premises at the site. In Topology #2, the WLCs and APs are deployed on-premises at a location such as a branch site, while the MSEs and PI are deployed remotely at a central site. In Topology #3 only the APs are deployed on-premises at a location such as a branch site, while most of the wireless infrastructure—consisting of WLCs, MSEs, and PI—are deployed centrally. Note that the network administrator may not necessarily have a choice in terms of the positioning of the WLCs in relation to the APs. In other words, the infrastructure may be an existing WLAN deployment which now requires CAS (location services) and/or CMX Services.

The benefits and disadvantages of each topology from a CAS and CMX perspective are summarized in Table 4-1.

Table 4-1 Benefits and Disadvantages of Deployment Topologies

Topology	Benefits	Disadvantages
#1: On-Premises WLCs, MSEs, and PI	<ul style="list-style-type: none"> <li>No additional WAN bandwidth requirements for transporting client RSSI information from APs to WLCs, and from WLCs to MSEs for processing location information.</li> <li>Potentially increased scale of the overall CMX deployment since each site has one or more MSEs.</li> </ul>	<ul style="list-style-type: none"> <li>Potentially increased cost of having to deploy and maintain WLCs and MSEs at each site.</li> <li>No ability to view location and analytics data across multiple sites from a central set of MSEs.</li> </ul>

**Table 4-1** *Benefits and Disadvantages of Deployment Topologies*

#2: On-Premises WLCs, Remote MSEs and PI	<ul style="list-style-type: none"> <li>• Potentially reduced cost of not having to deploy and maintain MSEs at each site.</li> <li>• The ability to view location and analytics data across multiple sites from a central set of MSEs.</li> </ul>	<ul style="list-style-type: none"> <li>• Potentially increased WAN bandwidth required to transport client RSSI information from the WLCs to a central set of MSEs for processing location information.</li> <li>• Potential scale limitations of the overall CMX deployment since all sites rely on a single set of MSEs.</li> </ul>
#3: Remote WLCs, MSEs, and PI	<ul style="list-style-type: none"> <li>• Potentially reduced cost of not having to deploy and maintain WLCs and MSEs at each site.</li> <li>• The ability to view location and analytics data across multiple sites from a central set of MSEs.</li> </ul>	<ul style="list-style-type: none"> <li>• Potentially increased WAN bandwidth required to transport client RSSI information from the APs to a central set of WLCs for processing location information.</li> <li>• Potential scale limitations of the overall CMX deployment since all sites rely on a single set of MSEs.</li> </ul>

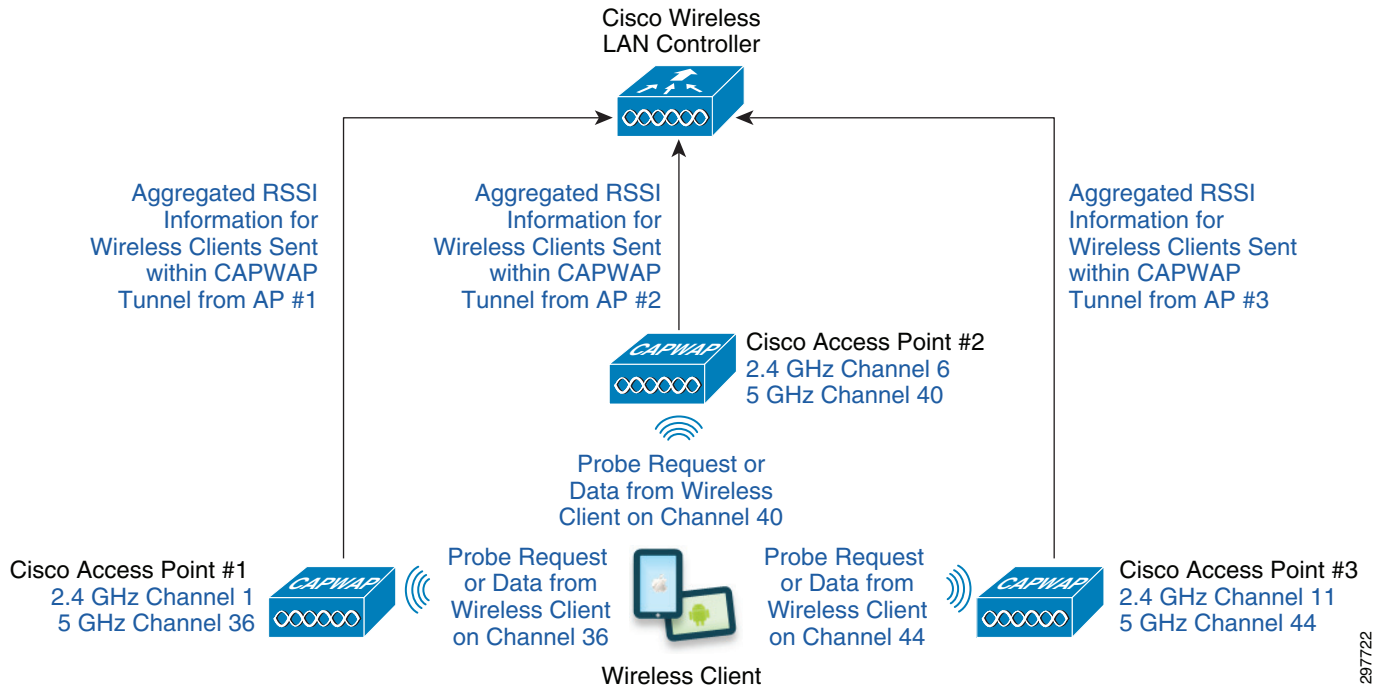
The following sections further discuss some of the considerations resulting from the deployment topologies.

## WAN Bandwidth Utilization

For deployments with limited amounts of bandwidth between the different sites where the MSEs, WLCs, and APs are deployed, the network administrator must ensure that sufficient bandwidth is provisioned to meet the requirements for existing applications at the site as well as the requirements for CAS (location services) and CMX services.

For CAS (location services), bandwidth is required to transmit aggregated RSSI data for each mobile device within the CAPWAP tunnel from each AP to the WLC, as shown in [Figure 4-2](#).

Figure 4-2 RSSI Information Sent within the CAPWAP Tunnel from APs to the WLC



The amount of bandwidth required for transmitting RSSI data between the APs and WLC depends upon multiple factors, including:

- The number of mobile devices at the site. Each mobile device either periodically generate Probe Requests or generate traffic which the FastLocate feature uses to calculate RSSI information. Hence the more mobile devices at the site, generally the more Probe Requests or traffic seen by FastLocate for a given time interval and the more bandwidth needed to accommodate the RSSI information within the CAPWAP tunnel.
- The frequency of packets generated by each mobile device which are used for location determination:
  - If Probe Request RSSI is implemented for CAS, then only Probe Requests from the wireless device are used by the AP to collect and send RSSI information to the WLC within the CAPWAP tunnel. Although the frequency of Probe Requests can vary from under a second to five minutes, a realistic number often used for the frequency of Probe Requests from mobile devices is 30 seconds.
  - If the FastLocate feature is implemented for CAS, then all packets from the wireless device are used by the WSM module within the AP to collect and send RSSI information to the WLC within the CAPWAP tunnel. With the FastLocate feature enabled, RSSI data can be collected from wireless clients as frequently as approximately every 4-6 seconds, depending upon whether CleanAir is implemented, the channels which the WSM scans, and whether a particular wireless client is transmitting when the WCM scans the particular channel on which the wireless client is operating. This could potentially increase the amount of RSSI traffic by 5-7 times over Probe Request RSSI. Alternatively, if wireless clients only periodically transmit packets such that FastLocate relies mostly on the BAR feature to update unresponsive clients, then there may be little to no increase in the amount of RSSI traffic over Probe Request RSSI.

- The number of APs deployed at the site which hear and report information on a given mobile device. If multiple APs hear either a Probe Request or data packets when using FastLocate, each AP collects RSSI information for the mobile device and send that information to the WLC. Therefore the more APs which hear the mobile device, the more bandwidth may be required to accommodate the RSSI information to the WLC within each of the CAPWAP tunnels.

The network administrator should note that Cisco APs automatically aggregate Probe Request information from mobile devices and forward them within CAPWAP messages to the WLC, regardless of whether or not an MSE is deployed within the infrastructure for CAS (location services) or CMX services. The aggregation timer for transmission of aggregated Probe Request information is, by default, set to 500 milliseconds.

**Note**

---

Probe Requests could come in at a rate high enough to fill the buffer used to hold the aggregated data before the aggregation timer expires. This could be due to a WLAN deployment with lots of mobile devices, each probing infrequently, or a WLAN deployment with fewer mobile devices, each probing frequently. In either case when the buffer has filled, the aggregated Probe Request information is sent immediately and the aggregation timer reset.

---

With the FastLocate feature enabled, APs with WSM modules collect RSSI values based on all data packets for all mobile devices heard while the radio within the WSM dwells on a particular channel. The Data RSSI values are again aggregated and forwarded to the WLC within CAPWAP messages. The aggregation timer is, by default, set to 500 milliseconds.

For Topologies #1 and #2 shown in [Figure 4-1](#), since the WLCs and APs are both on-premises, it is typically assumed there is sufficient bandwidth on the LAN to accommodate the RSSI information. Often the APs are deployed in an overlay network (Local Mode) in which all wireless traffic is backhauled to the WLC before being terminated on the LAN. Hence the amount of traffic due to RSSI information within the CAPWAP tunnel is also typically small compared to the actual WLAN traffic encapsulated within the CAPWAP tunnel.

For Topology #3 however, since the WLCs are remote from the APs, the network administrator must ensure sufficient bandwidth on the WAN to accommodate the RSSI information. Topology #3 is indicative of a typical FlexConnect WLAN deployment within a small branch in which wireless traffic is often terminated locally at the AP.

**Note**

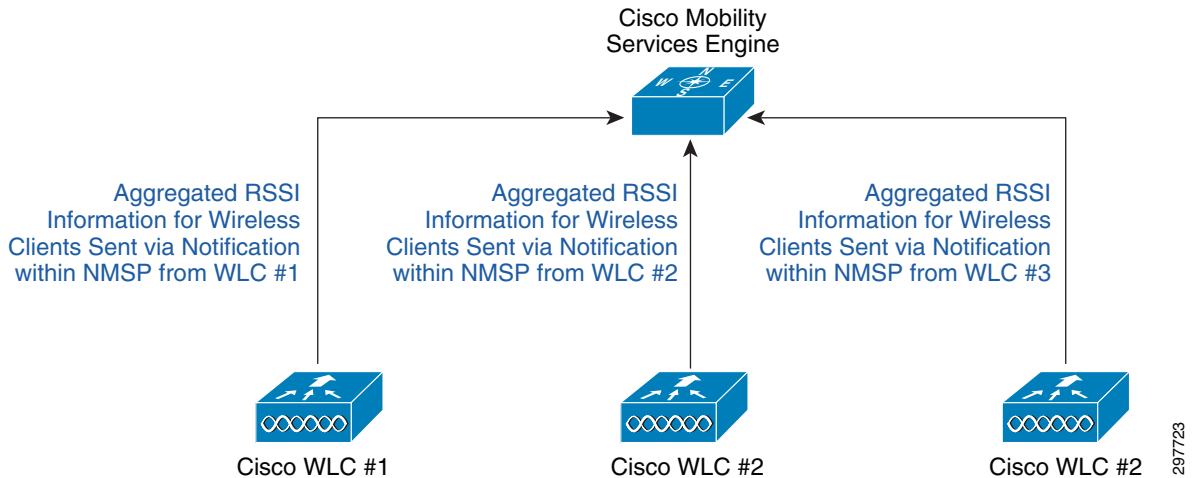
---

The CAPWAP control channel between the APs and WLCs has multiple functions besides transmitting aggregated RSSI information. Hence additional bandwidth overhead is incurred when deploying a FlexConnect WLAN design. These details are not covered within this design guide.

---

For CAS (location services), bandwidth is also required to transmit aggregated RSSI information, within NMSP Measurement Notification messages, from each WLC to the MSE, as shown in [Figure 4-3](#).

**Figure 4-3** Aggregated RSSI Information Sent within the CAPWAP Tunnel from APs to the WLC



The amount of bandwidth required for transmitting RSSI data between the WLCs and the MSE depends on the same factors as discussed in the previous section since the RSSI information from the APs is aggregated by the WLCs before being sent to the MSE. The amount of required bandwidth is also dependent on the following additional factors:

- The number of WLCs deployed. The overall deployment may consist of multiple sites, each with its own WLC, or a single site with multiple WLCs. Each WLC transmits all NMSP messages (containing information for all of its APs) to all MSEs associated with it. RSSI Notification Messages sent within NMSP contain the aggregated RSSI information. These are sent every two seconds by default.
- The number of MSEs deployed. Again, each WLC sends all NMSP messages to all MSEs associated with it (regardless of maps). Hence if a given WLC has two MSEs associated with it, the amount of NMSP traffic is effectively doubled.

For Topologies #1 and #3 shown in Figure 4-1, since the WLCs and MSEs are both on-premises, it is typically assumed there is sufficient bandwidth on the LAN to accommodate the RSSI information.

For Topology #2 however, since the MSE is remote from the WLCs and APs, the network administrator must ensure sufficient bandwidth is provisioned within the WAN to accommodate the RSSI information.



**Note**

The NMSP protocol between the WLC and MSE has multiple functions besides transmitting aggregated RSSI information, such as transmitting wIPS information. Hence additional WAN bandwidth overhead may be incurred when deploying a design in which the WLCs are remote from the MSEs. These details are not covered within this design guide.

For CMX services, additional WAN bandwidth may be required to support guest traffic if CMX Visitor Connect is deployed at the site and if the guest traffic is sent back to a central location (auto-anchored) before being sent to the Internet. The amount of WAN bandwidth required depends on the number of guests simultaneously supported at the site and also varies highly based on whether the guest is simply browsing a web site or downloading streaming video. Per user rate-limiting can be applied to the B2C guest WLAN at the WLC to limit the amount of bandwidth that each guest can consume to partially alleviate this issue.

**Note**

The guest WLAN traffic can be sent to the Internet directly from the site. Future versions of the CMX design guide may address these designs.

Additional WAN bandwidth may be required to support push notifications and/or application specific information to the mobile device. This is required if the mobile device is again associated to the guest WLAN at the site and the client is configured with a CMX mobile application and if the guest traffic is auto-anchored back to the central location before being sent to the Internet. Future versions of the CMX design guide will discuss the Cisco CMX Mobile Application Server & CMX SDK.

It is assumed that end-users may also access location and presence analytics data both from within branches and from within campus locations. Hence bandwidth may also be utilized when end-users access analytics data from within a branch and the MSE is located within a centralized campus location.

## MSE Scalability

The scalability of the Cisco MSE, in terms of licensing, is shown in [Table 4-2](#). As of MSE software release 7.4 and above, licensing is based the number of APs supported.

**Table 4-2** *Mobility Services Engine Scalability*

<b>Platform</b>	<b>Location Services Licensing</b>	<b>Advanced Location Services Licensing</b>	<b>wIPS Licensing (Monitor Mode or Enhanced Local Mode)</b>
Cisco 3355 MSE Appliance	Up to 2,500 APs	Up to 2,500 APs	Up to 5,000 APs
Cisco MSE Virtual Appliance (High-end Server)	Up to 5,000 APs	Up to 5,000 APs	Up to 10,000 APs
Cisco MSE Virtual Appliance (Standard Server)	Up to 2,500 APs	Up to 2,500 APs	Up to 5,000 APs
Cisco MSE Virtual Appliance (Basic Server)	Up to 200 APs	Does not support Advanced Location (CMX) Services	Up to 2,000 APs

As of MSE release 7.5 and higher, there is no enforcement regarding the number of end-devices on the MSE. However there is a hard limit of 25,000 tracked end-devices on the Cisco 3355 MSE appliance and 50,000 tracked end-devices on the on the high-end virtual MSE server.

Since location calculations are CPU intensive, the MSE also scales based on the number of location calculations it has to perform. A high-end virtual MSE server can handle approximately 90,000 calculations per minute (approximately 1,500 calculations per second). A new location calculation is performed by the MSE for a given wireless client when the following conditions occur:

- New RSSI data is seen by the MSE for the wireless client, which can occur for the following reasons:
  - With Probe Request RSSI—The wireless client sent Probe Requests.



- With the FastLocate feature enabled —The wireless client sent packets which were seen by WSM modules within the APs during the dwell time on the particular channel the wireless client was transmitting on.
- The signal strength is more than ~5 dBm from the last time RSSI information was received from the wireless client (i.e., the last time the wireless client was seen). Note that this is a configurable parameter. Information about setting this parameter is provided in [Chapter 25, “Configuring the Mobility Services Engine for CMX.”](#)

In dense AP deployments, if more than five APs see the wireless client, the MSE only uses the top five (the strongest RSSI values) in calculating the X,Y coordinates of the wireless client. However the MSE expends additional CPU cycles sorting and discarding the excess data points.

A very conservative estimate of the number of devices which can be supported by a high-end virtual MSE server would be to assume that updated RSSI information is received from every wireless client every 5-6 seconds (~10 updates per minute). This corresponds to the example of using FastLocate with CleanAir discussed in [Probe Request RSSI versus FastLocate in Chapter 3, “CMX Solution Components.”](#) Using a maximum 90,000 calculations per minute for a high-end virtual MSE, yields a maximum number of clients for the MSE as follows:

$90,000 \text{ updates per minute} / 10 \text{ updates per wireless client per minute} = 9,000 \text{ wireless clients}$

Hence a very conservative estimate of the number of devices which can be supported by a high-end virtual MSE server is 9,000 wireless devices seen by WLC and tracked by MSE.

As mentioned in [Probe Request RSSI versus FastLocate in Chapter 3, “CMX Solution Components,”](#) this would require each wireless client to be transmitting packets during the time the WCMs dwell upon the channel which each wireless client is operating on during every scan cycle. Assuming the RSSI information from the 9,000 wireless clients is averaged out—meaning that for every second, approximately 1/6th of the 9,000 wireless clients are heard from—results in a movement percentage of approximately 16.67%.

A more realistic estimate of the number of devices which can be supported by a high-end virtual MSE server would be to assume that updated RSSI information is received from every wireless client every 10-30 seconds (~2-6 updates per minute). Again, using a maximum 90,000 calculations per minute for a high-end virtual MSE yields a maximum of clients for the MSE as follows:

$90,000 \text{ updates per minute} / 6 \text{ updates per wireless client per minute} = 15,000 \text{ wireless clients}$

$90,000 \text{ updates per minute} / 2 \text{ updates per wireless client per minute} = 45,000 \text{ wireless clients}$

Hence a more realistic estimate of the number of devices which can be supported by a high-end virtual MSE server is from 15,000-45,000 wireless devices seen by WLC and tracked by MSE. Again, assuming the movement of the wireless clients is averaged out, results in a movement percentage from 3.33% - 10%.



#### Note

The example above highlights an alternative method of estimating scale of the MSE based on movement percentage of wireless devices. For example, for a given a movement percentage of 3%, an estimate of the number of wireless clients a large MSE virtual server can handle is approximately 50,000 wireless devices.

An additional technique which can help scale is to not overprovision cores on the virtual machine which hosts the MSE. In other words the number of virtual CPUs (vCPUs) should equal the number of physical cores. This gives the MSE more CPU time to calculate locations at high scale. Also, a hard disk which supports 1,600 input/output operations per second (IOPS) is recommended. The minimum of 250 IOPS can be used, but may result in slow updates to client location using the MSE API and sluggish response when accessing the MSE graphical user interface (GUI).



Further scaling of the MSE can be accomplished by splitting out services, which may be necessary for large scale WLCs which see more than 50,000 devices and when CMX Analytics and CAS are deployed. The following table shows possible ways of additional MSE scaling by splitting out services.

**Table 4-3** *MSE Scaling by Splitting Out Services*

Number of MSEs	Services Running on Each MSE	Scale	Devices	Caveats
1	CAS Analytics wIPS Connect/Engage Mobile Concierge	Demo purposes—up to 1,000 devices	1,000	Recommended to run aWIPS on a separate MSE, since aWIPS requires UTC time zone. Otherwise Analytics information will have the incorrect time zone.
1	CAS Analytics Connect/Engage Mobile Concierge	Demo purposes—up to 1,000 devices without wIPS	1,000	No aWIPS. Good for small deployments or demonstrations.
2	1 CAS 1 Analytics	Highest scale for CAS to run on a separate box. Scale to 50,000 devices with 3% movement rate	10,000–50,000 per WLC/MSE	No aWIPS, CMX SDK, or guest services such as CMX Visitor Connect. MSE REST API and Analytics only.
4 + 1 CMX Mobile App Server	1 CAS 1 Analytics 1 wIPS 1 Connect/Engage & Mobile Concierge 1 CMX Mobile App Server	Highest overall scale by separating all services into individual chassis.  Use this architecture when there are over 10,000 clients and you need SDK + Analytics + Visitor Connect	10,000–50,000 per WLC/MSE	Splitting out Analytics save ~ 15% CPU when extracting analytics data from raw db.  Splitting out Connect/Engage save ~15% CPU for rendering Guest Login web pages under load.  Adding CMX Mobile App Server decreases latency for location API calls.

## Campus and Branch Designs

Applying the basic high-level topologies (as shown in [Figure 4-1](#)) to traditional enterprise campus and branch designs which may require location and CMX services results in the customer deployment models shown in [Table 4-4](#).

T

**Table 4-4** *Campus and Branch Deployment Models*

Deployment Model	Infrastructure	Environment
Single Campus (or Large Branch) Deployment Model	On-Premises APs, WLCs, MSEs, and PI	This type of deployment applies to single-site campus networks, such as a standalone hospital or a single-campus college/university. This type of deployment also applies to single-site large branch network, such as a large retail store.
Multiple Campuses (or Large Branches) Deployment Model	On-Premises APs and WLCs; Remote MSEs and PI	This type of deployment applies to networks that support multiple large campus sites, such as regional hospitals and universities with multiple campuses. This type of deployment also applies to networks that support multiple large-sized branches, such as large retail stores connected to a campus site.
Small Branch Deployment Model	On-premises APs; Remote WLCs, MSEs, and PI	This type of deployment applies to networks that support multiple small-sized remote branches, such as small retail stores or branches of financial institutions which connect to a campus site.

This version of the CMX design guide highlights two of the deployment models—the Single Campus (or Large Branch) Deployment Model and the Small Branch deployment model—in relation to the CMX Location Analytics, CMX Presence Analytics, and CMX Visitor Connect use cases.

## Single Campus (or Large Branch) Deployment Model

From the perspective of CAS and CMX services, the difference between a single campus and a single large branch is the really just the size of network infrastructure within the site. Campus networks tend to have multiple buildings, requiring a traditional three-tiered (core, distribution, access) switch infrastructure and a more modular design (Internet Edge Module, Services Module, Data Center Module, etc.) for separation of function and scalability. Large branch networks tend to be much simpler, requiring a two-tiered (distribution and access) switch infrastructure and often no modularity within the design. Both campus and large branch network have the APs, WLCs, MSEs, and PI in the same physical location relative to each other. Hence these have been combined into a Single Campus Deployment Model in the following discussion for simplicity.

The Single Campus Deployment Model is the topology used for the CMX Location Analytics use cases provided in [Chapter 7, “CMX Use Case Stories.”](#)

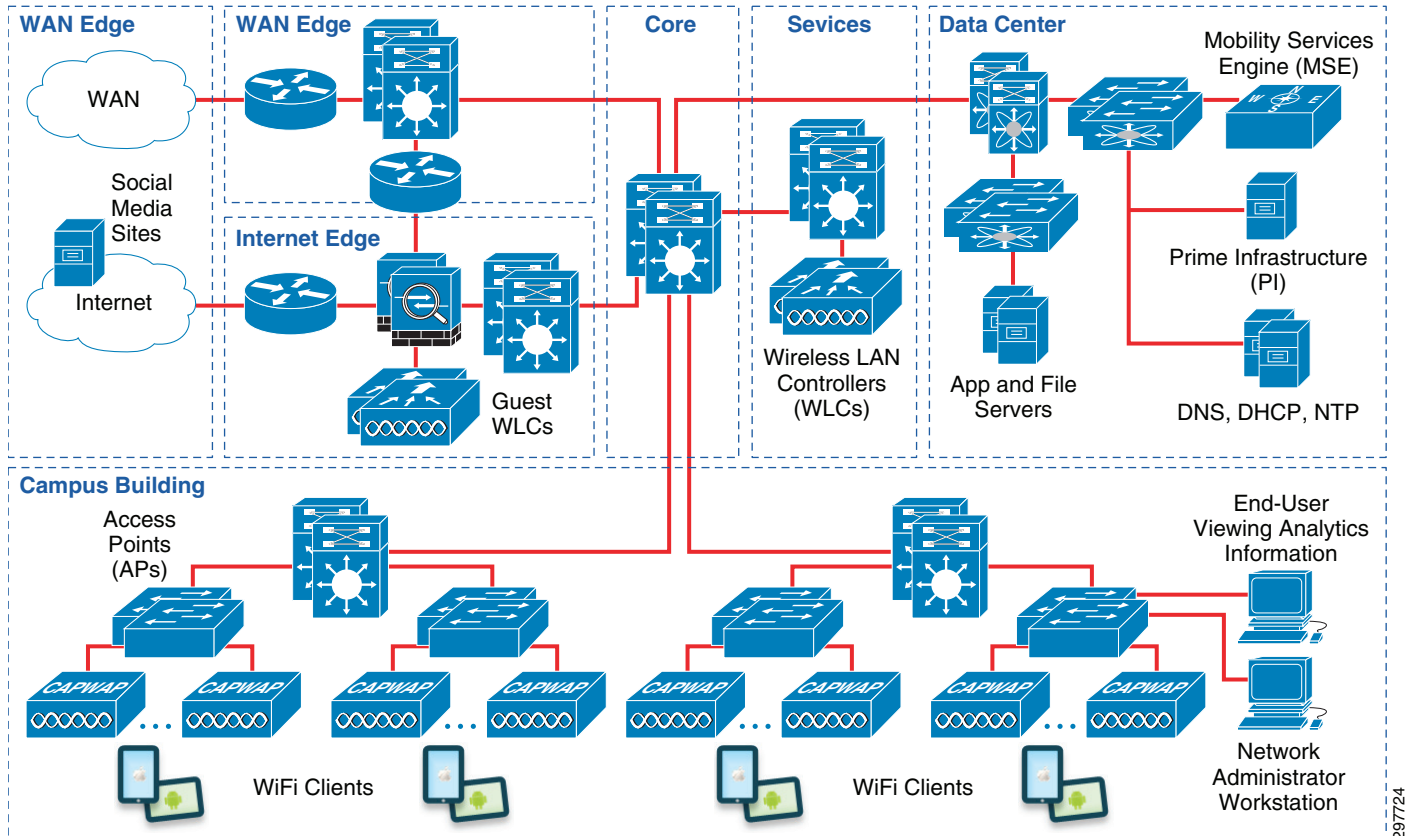


### Note

A campus deployment is shown as the example in this section for simplicity.

Figure 4-4 shows the placement of the wireless infrastructure components (APs, WLCs, MSE, and PI) within the campus for the Single Campus Deployment Model.

Figure 4-4 Single Campus Deployment Model



This deployment model is the classic centralized (Local Mode) design applied to a campus site which has sufficient infrastructure to support a three-tiered network design. Cisco second generation 3700, 3600, 2700, or 2600 APs deployed within a campus are controlled by Cisco 5500 Series WLCs deployed locally within a separate services module of the campus.

Wireless data traffic is terminated centrally (Local Mode) on the wireless controllers, with the exception of the B2C guest WLAN/SSID which is discussed in [B2C Guest Access for CMX Visitor Connect](#). One or more MSEs (either 3355 standalone appliances or virtual machines running on Cisco UCS infrastructure) are also deployed within the campus data center to support the Context Aware Service (location), CMX Analytics, and CMX Visitor Connect. Cisco Prime Infrastructure running on a virtual machine also within the data center is used to import, configure, and synchronize floor maps to the MSE as well as enable the MSE services.

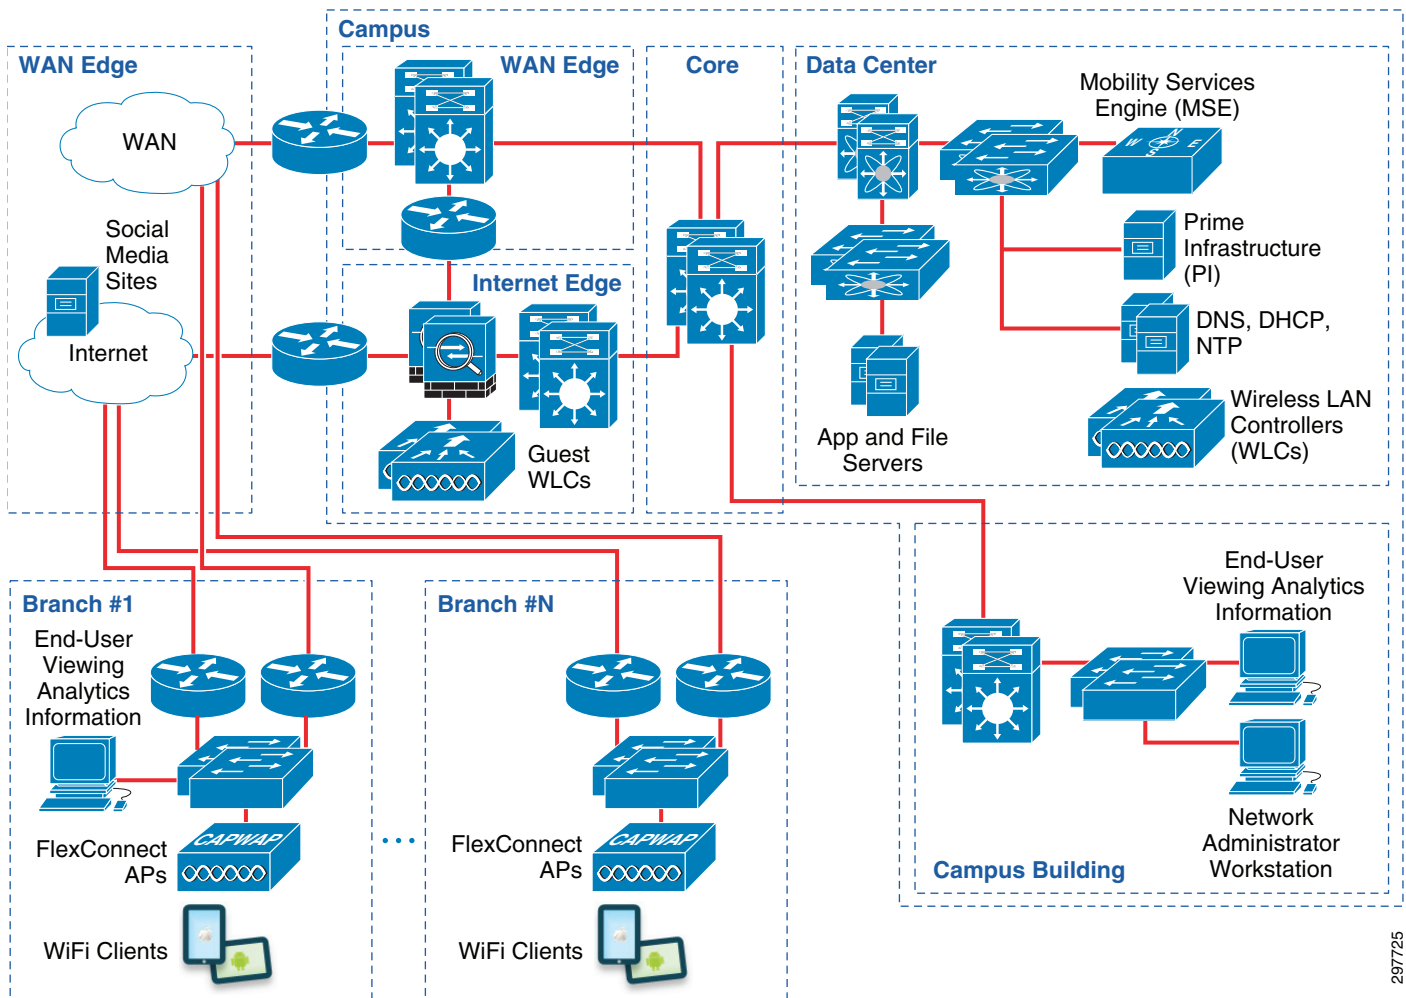
Since the APs, WLCs, and MSE are all on-premises, no WAN bandwidth is utilized in transporting RSSI information from APs to the WLCs and from the WLCs to the MSE. This is one of the advantages of this design. However a disadvantage of this design is that a system-wide view as it relates to Cisco CleanAir, WIPS, Context Aware Services, or CMX cannot be achieved if each Campus location runs its own set of MSEs.

## Small Branch Deployment Model

The Small Branch deployment model is the topology used for the CMX Presence Analytics use case and the CMX Visitor Connect use case that is provided in [Chapter 7, “CMX Use Case Stories.”](#)

[Figure 4-5](#) shows the placement of the wireless infrastructure components (APs, WLCs, MSE, and PI) both within the branch and the campus, for the Small Branch Deployment Model.

**Figure 4-5** Small Branch Deployment Model



This deployment model is suited for customer deployments in which a large number of branches, each of which only requires one or two APs, need to be supported. Cisco second generation 3700, 2700, 3600, or 2600 Series APs within each branch are configured to operate in FlexConnect mode. APs deployed within the multiple small branches are controlled by Cisco Flex 7500 Series WLCs located within the data center of a remote campus.

Wireless data traffic is terminated directly on the APs, with the exception of the B2C guest WLAN/SSID, which is discussed in [B2C Guest Access for CMX Visitor Connect](#). One or more MSEs (either 3355 standalone appliances or virtual machines running on Cisco UCS infrastructure) are also deployed within the campus data center to support the Context Aware Service (location), CMX Analytics, and

CMX Visitor Connect. Cisco Prime Infrastructure running on a virtual machine also within the data center is used to import, configure, and synchronize floor maps to the MSE as well as enable the MSE services.

The bandwidth utilization of transporting RSSI information from each AP to the WLC is one of the disadvantages of this design. However the advantage of this design is the lower cost of supporting one or more centralized MSEs within a campus location which service all branches versus deploying one or more MSEs at each branch. Further, a system-wide view as it relates to Cisco CleanAir, wIPS, Context Aware Services, or CMX can be achieved if the MSE is servicing WLCs throughout the network.

## B2C Guest Access for CMX Visitor Connect

For this version of the Cisco CMX design guide, guest wireless Internet access is provided by auto-anchoring the traffic from a B2C guest WLAN/SSID back through either the Flex 7500 Series WLCs (for the Small Branch Deployment model) or the 5500 Series WLCs (for the Single Campus Deployment model) and terminating the traffic on a DMZ segment off of the dedicated guest 5508 WLC within the Internet Edge of the campus. The B2C guest WLAN/SSID, used for CMX Visitor Connect, is configured for open authentication with Web Passthrough on both the internal (foreign) WLC and the guest (anchor) WLC. The MSE is configured as the external web portal within the WLCs for web redirection. Details for configuring the WLCs are provided in [Chapter 23, “Configuring Cisco Wireless LAN Controllers.”](#)





## CMX Security Considerations

---

September 4, 2014

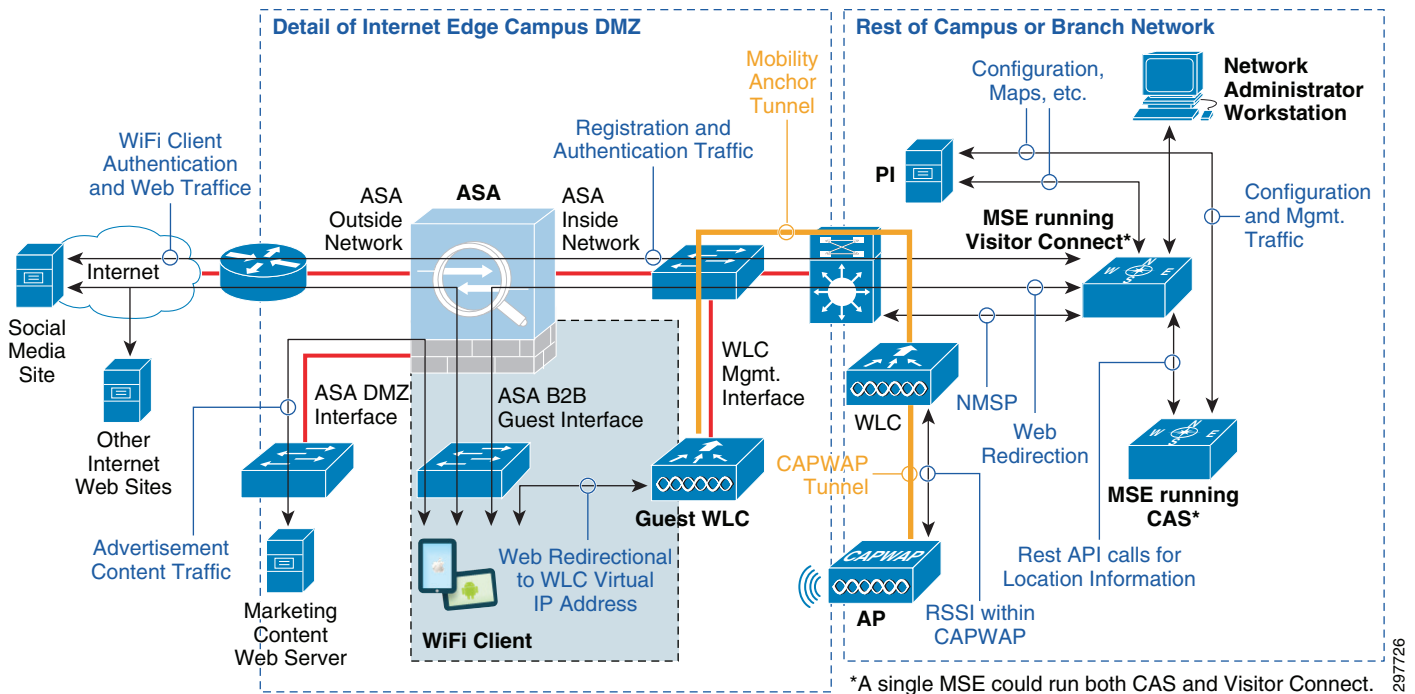
This chapter focuses on traffic isolation for guest wireless access as part of CMX Visitor Connect. Additionally, it discusses Role-Based Access Control (RBAC) for the Mobility Services Engine (MSE) as well as the CMX Connect & Engage service.

### Traffic Isolation for CMX Visitor Connect

With CMX Visitor Connect, guests must be allowed to authenticate to the network using credentials from social media sites, such as Facebook, LinkedIn, and Google+. This involves the use of a variation of the OAuth protocol along with the MSE. Hence the MSE which runs CMX Visitor Connect must be reachable by guests. However all other guest traffic (other than DHCP and DNS) should be isolated from the rest of the corporate network. Guests are only allowed to access the Internet in this design.

[Figure 5-1](#) shows at a high-level the hardware and various flows involved in CMX Visitor Connect.

Figure 5-1 Hardware and Information Flows for CMX Visitor Connect



With this design, traffic isolation is primarily achieved via policy configured on the ASA firewall operating as a Layer 3 firewall within the Internet Edge DMZ. Note that virtualization of the network infrastructure is not utilized in this design to isolate guest traffic.

The following interfaces are configured on the ASA firewall for this design example:

- ASA Inside Interface—This has the highest security level. The IP addressing is assumed to be private and part of the enterprise organization for this design. An example policy for the ASA Inside Interface is as follows:
  - Allow all traffic from all devices initiated from any subnet reachable via the ASA Inside Interface to all devices which are reachable via any lower security level interfaces (ASA DMZ Interface, the ASA B2C Guest Interface, and the ASA Outside Interface). The ASA firewall should allow return traffic from any lower security interface back to a higher security interface, as long as the session was initiated from a device on the higher security interface. This allows the necessary connections from the MSE to social media sites for authentication for CMX Visitor Connect.
  - It is assumed for this design that NAT is not used between the ASA Inside Interface and the ASA DMZ Interface or the ASA B2C Guest Interface.
  - It is assumed for this design that NAT is used between the ASA Inside Interface and the ASA Outside Interface since the ASA Outside Interface has a publicly routable IP address.
- ASA Outside Interface—This has the lowest security level. The IP addressing is assumed to be publicly routable in this design. An example policy for the ASA Outside Interface is as follows:
  - Block all traffic from all devices initiated from any subnet reachable via the ASA Outside Interface to all devices which are reachable via any higher security level interfaces (ASA DMZ Interface, the B2C Guest Interface, and the ASA Outside Interface). This effectively blocks any sessions initiated from the Internet from reaching any resources within the corporation, while still allowing return traffic from any sessions initiated from within the corporation out to the Internet.



- It is assumed that NAT is used between the ASA Outside Interface and the ASA DMZ Interface, the ASA B2C Guest Interface, and the ASA Inside Interface for this design.
- ASA B2C Guest Interface—This has a security level below the ASA Inside Interface and the ASA DMZ Interface, but above the ASA Outside Interface for this design. The IP addressing is assumed to be private and part of the enterprise organization. An example policy for the ASA B2C Guest Interface is as follows:
  - Allow inbound DNS from any device with a source IP address on the ASA B2C Guest Interface subnet to the corporate DNS server located on a subnet available via the ASA Inside Interface.
  - Allow inbound traffic from any device with a source IP address on the ASA B2C Guest Interface subnet destined to TCP port 8083 of the IP address of the MSE which is running CMX Visitor Connect. This is necessary for the redirection of the B2C guest web browser to CMX Visitor Connect running on the MSE during the authentication via social media sites process.
  - Allow inbound traffic from any device with a source IP address on the ASA B2C Guest Interface subnet destined to TCP port 80 (HTTP) of the IP address of the Marketing Content Server which is sitting on the ASA DMZ Interface. This is necessary for allowing the B2C guest browser to access any advertisement content on the Marketing Content Server without allowing the B2C guest device onto the inside of the corporate network. The ability to display advertisement content on the B2C guest device is an optional step in the three-step Visitor Connect process.
  - DHCP should be configured to be relayed by the dedicated Guest WLC to a DHCP server located on the inside of the corporate network. In this configuration, the Mgmt. Interface of the Guest WLC is the source IP address used by the DHCP relay function. Hence the Mgmt. Interface of the Guest WLC needs to have an IP address which is part of the corporate address space reachable via the ASA Inside Interface. In this configuration, the ASA B2C Guest Interface should not need to include an access entry allowing inbound DHCP traffic to the ASA Inside Interface.
  - By default allow all traffic from all devices initiated from any IP address on the ASA B2C Guest Interface subnet to all devices which are reachable via any lower security level interfaces (ASA Outside Interface is the only lower security interface). This allows B2C guests to access the Internet and perform authentication to the social media sites as part of CMX Visitor Connect.
  - It is assumed that NAT is not used between the ASA B2C Guest Interface and the ASA DMZ Interface, or the ASA Inside Interface.
  - It is assumed that NAT is used between the ASA B2C Guest Interface and the ASA Outside Interface since the ASA Outside Interface has a publicly routable IP address.
- ASA DMZ Interface—This has a security level below the ASA Inside Interface, but above the ASA B2C Guest Interface and the ASA Outside Interface. The IP addressing is assumed to be private and part of the Enterprise organization. The ASA DMZ Interface is assumed to house the Marketing Content Web Server. Note that in an actual deployment, the marketing content may be deployed in the cloud and accessible from the Internet, rather than onsite on a DMZ server. Hence this may be optional. An example policy for the ASA DMZ Interface is as follows:
  - By default block all traffic initiated from any device with a source IP address on the ASA DMZ Interface subnet to any other interface. The Marketing Content Web Server is supposed to respond to HTTP requests for marketing content, not generate any requests on its own. This should still allow the Marketing Content Web Server to respond to HTTP requests for content which are initiated from devices on the ASA B2C Guest Interface subnet.
  - It is assumed that NAT is not used between the ASA DMZ Interface and the ASA Inside Interface or the ASA B2C Guest Interface.
  - It is assumed that NAT is used between the ASA DMZ Interface and the ASA Outside Interface since the ASA Outside Interface has a publicly routable IP address.

Note that locking down the ASA firewall such that guests using CMX Visitor Connect can only access TCP port 8083 is a necessary step for securing the overall CMX deployment. The MSE may be running multiple services, including the Context Aware Service (CAS) and CMX Analytics. Hence tight network access control along with role-based access control to the MSE is recommended to minimize chances of unauthorized access to the MSE.

## Role-Based Access Control on the MSE

The MSE itself has its own role-based access control (RBAC) separate from the CMX Connect & Engage service. Role-based access control on the MSE controls access to administrative functions on the MSE itself, as well as access to CMX Analytics (both Location and Presence). For role-based access control of the CMX Connect & Engage service, see [Role-Based Access Control for the CMX Connect & Engage Service](#).

Services such as CMX Analytics are intended to be utilized by non-IT personnel. For example, in a retail deployment, store operations managers may need to access the MSE running CMX Analytics to view the dashboard, run reports, or run custom analysis. However IT personnel should be the only ones allowed to modify the configuration of the services running on the MSE and shutdown or restart the MSE. Hence the implementation of role-based access control is considered an essential security measure to mitigate the chances of any accidental or malicious disruption of the services (CMX and/or CAS) provided by the MSE.

Role-based access control (RBAC) consists of configuring one or more Groups on the MSE with one of the following three access-control privileges—Read Access, Write Access, or Full Access. Individual Users are then created and assigned to one of the Groups. All Groups and Users reside on the local database within the MSE. There is no integration with an external data store, such as an LDAP database or external RADIUS server.

As of MSE version 8.0, it is not recommended to utilize Read Access groups on the MSE. CMX Presence Analytics currently does not participate in Role-Based Access Control (RBAC) on the MSE. Hence any userid which is part of a Read Access group is also able to add/modify/delete CMX Presence Analytics configuration. Note that the use of only Write Access and Full Access groups means that any non-IT personnel who require access to the CMX Analytics dashboard, analytics tab, or reports may also have access to add/modify/delete CMX Analytics (location and presence) configuration. One way to mitigate some of this risk is for IT personnel to download CMX Analytics Reports and email them to non-IT personnel at regular intervals. Alternatively, the list of non-IT personnel—such as store operations managers, marketing executives, etc.—who have direct access to the MSE for CMX Analytics should be kept tightly controlled.

The MSE does support the ability to enforce the choice of a strong password, meaning a minimum length of password which includes capital letters and special characters. The MSE does not provide any mechanism for the end user to change their password after a certain period of time or to change their password at all. Hence the MSE administrator must be responsible for all user accounts and should always choose a strong password when creating user accounts. The MSE does not support the ability to disable a user password after a number of unsuccessful attempts. Hence there is limited ability to guard against unauthorized access to the MSE using a dictionary attack. This makes it all the more critical that the network administrator choose a strong password, especially for user accounts which belong to groups which have Write Access or Full Access.

The MSE supports the ability to monitor active sessions. The MSE administrator can view active sessions by selecting the Active Sessions link under the System topic from the MSE Dashboard page. This displays the Active Sessions page, as shown in [Figure 5-2](#).

**Figure 5-2** Example of the Active Sessions Page

Session Identifier	Access from (Host)	Username	Time Started	Last Access	Idle (secs)
1392	10.230.1.102	admin	Jun-12-2014 02:51:22 AM	Jun-24-2014 05:09:24 AM	3

The length of time that a user's session is idle before being logged out of the MSE is set for 30 minutes. This setting is currently viewable under the Advanced Parameters settings, but is not configurable. An example is shown in Figure 5-3.

**Figure 5-3** Session Timeout Parameter

**Advanced Parameters**

Number of days to keep events  
 1 - 365 days

Session Timeout  
 minutes

**Advanced Commands**

The MSE administrator can also view logs to monitor activity of users on the MSE. The MSE administrator can view the logs by selecting the Audit Logs link under the Status topic from the MSE Dashboard page. This displays the Audit Logs page, as shown in Figure 5-4.

Figure 5-4 Example of the Audit Logs Page

User Name	Operation	Operation Status	Module	Invocation Time
-1	Update Track Group corresponding to Rest API Notification Subscription, name: admin/bbx-event	SUCCESS	ADMIN	Jun-04-2014 09:47 AM
admin	MSE services modified. Context Aware Service ENABLED. WIPS DISABLED. Mobile Concierge Service DISABLED. CMX Analytics ENABLED. CMX Browser Engage ENABLED. HTTP Proxy Service DISABLED.	SUCCESS	ADMIN	Jun-04-2014 09:43 AM
admin	MSE services modified. Context Aware Service ENABLED. WIPS DISABLED. Mobile Concierge Service DISABLED. CMX Analytics DISABLED. CMX Browser Engage ENABLED. HTTP Proxy Service DISABLED.	SUCCESS	ADMIN	Jun-04-2014 09:43 AM
admin	MSE services modified. Context Aware Service ENABLED. WIPS DISABLED. Mobile Concierge Service DISABLED. CMX Analytics DISABLED. CMX Browser Engage ENABLED. HTTP Proxy Service DISABLED.	SUCCESS	ADMIN	Jun-04-2014 09:43 AM
admin	MSE services modified. Context Aware Service ENABLED. WIPS DISABLED. Mobile Concierge Service DISABLED. CMX Analytics DISABLED. CMX Browser Engage ENABLED. HTTP Proxy Service DISABLED.	SUCCESS	ADMIN	Jun-04-2014 09:43 AM
-1	Add Track Group corresponding to Rest API Notification Subscription, name: admin/bbx-event	SUCCESS	ADMIN	Jun-04-2014 09:36 AM

The audit logs provide a fairly comprehensive audit trail, showing the specific operation which was performed, the user who performed that operation, whether it was successful or not, and the time when the operation was performed.

## Role-Based Access Control for the CMX Connect & Engage Service

The CMX Connect & Engage service (which includes CMX Visitor Connect) has its own role-based access control (RBAC) separate from the MSE itself. For role-based access control on the MSE, see [Role-Based Access Control on the MSE](#).

RBAC on the CMX Connect & Engage service consists of configuring Roles which can perform one or more operations. Individual Users are then created and assigned to one of the Roles. As of MSE version 8.0, all Roles and Users reside on the local database within the CMX Connect & Engage service on the MSE. There is no integration with an external data store, such as an LDAP database or external RADIUS server.

CMX Connect & Engage provides very granular role-based access control (RBAC). Each role can be configured for each of the following 15 operations:

- Accounts—Allows members of the role to create, edit, and delete Accounts. Accounts are associated with different Campaigns and Banners.
- Banner Approver—Allows members of the role to approve Banners.
- Banners—Allows members of the role to create and edit Banners.
- Campaigns—Allows members of the role to create and edit Campaigns.
- Campaign Approver—Allows members of the role to approve Campaigns.
- CMX Mobile—Allows members of the role to access functions found under the Mobile App topic.
- Domain Setup—Allows members of the role to create, edit, and delete domains, which are found under the Settings topic.
- Floor Navigation—Allows members of the role to access Floor Navigation functions found under the Mobile App topic.

- Menu
- Point of Interest—Allows members of the role to create, edit, and delete points of interest.
- Reports
- Roles—Allows members of the role to create, edit, and delete Roles.
- Server Settings—Allows members of the role to access Server Settings functions found under the Settings topic.
- Users—Allows members of the role to create, edit, and delete Users. Users are associated to a particular role for role-based access control.
- Visitor Connect—Allows members of the role to create splash templates and configure social media connectors.

As of software version 8.0, the CMX Connect & Engage service does support the ability to enforce the choice of a strong password—meaning a minimum length of password which includes capital letters and special characters. The CMX Connect & Engage service does not provide any mechanism for the end user to change their password after a certain period of time or to change their password at all. Hence the CMX administrator must be responsible for all user accounts and should always choose a strong password when creating user accounts.

CMX Visitor Connect & Engage does not support the ability to disable a user password after a number of unsuccessful attempts. Hence there is limited ability to guard against unauthorized access to the CMX Connect & Engage service using a dictionary attack. This makes it all the more critical that the network administrator choose a strong password, especially for users who belong to roles such as the Super Admin, who have access to all CMX Connect & Engage operations.





## CMX Additional Considerations

---

September 4, 2014

This chapter highlights additional considerations when deploying a CMX solution. This includes how fast location information is updated and made available, considerations around specific mobile device platforms such as Apple iOS 8 devices and some Android devices, considerations around the use of 2.4 and 5 GHz frequency bands when deploying location services and CMX services, and finally considerations around the deployment of the FastLocate feature.

### Currency of Location Information

This section discusses how fast location information is processed by the Cisco Context Aware Service (CAS) and made available to devices.

The default WLC aggregation window of two seconds, discussed in [Cisco Context Aware Service \(CAS\) in Chapter 3, “CMX Solution Components,”](#) results in an MSE aggregation window of five seconds. Adding in the time required to calculate the location of the wireless client, send a push notification, and update the location database results in the updated location of the client being available in as much as approximately eight seconds after the client was heard via either Probe Requests or via data packets when using the FastLocate feature.

Adjusting the WLC aggregation window down to one second results in an MSE aggregation window of four seconds. Adding in the time required to calculate the location of the wireless client, send a push notification, and update the location database results in the updated location of the client being available in as much as approximately six seconds after the client was heard via either Probe Requests or via data packets when using the FastLocate feature.

Changing the default WLC aggregation window from two seconds to one second does have implications on the overall scalability of the MSE since it will need to process location information more rapidly. [MSE Scalability in Chapter 4, “CMX Deployment Models”](#) discusses the scalability of the MSE.

As discussed in [Probe Request RSSI versus FastLocate in Chapter 3, “CMX Solution Components,”](#) the frequency of Probe Requests from mobile devices is often non-deterministic and may vary from under a second to five minutes, depending upon various factors. Some mobile device platforms may allow an app developer the ability to generate Probe Requests, via the app, while others may not. The deployment of the FastLocate feature can result in RSSI information being seen from the mobile device in as little as every four to six seconds. However this is only if the mobile device generates packets for every scan cycle of the WSM. This may again require an app running on the mobile device to constantly generate traffic. For unresponsive clients, the Block Acknowledgement Request (BAR) feature of FastLocate may still result in the mobile device being seen, but only every 40 to 60 seconds.

It must be recognized by any app developer that the 6-8 second delay in the availability of location information discussed above must be added to the frequency by which the mobile device generates Probe Requests or generates packets when using the FastLocate feature to understand the total delay in the availability of location information regarding that particular mobile device. In addition, this does not include time required to process any push notifications via an application running on a server or an app running on the mobile device itself. Nor does it include transmission delays in sending the push notification to an application running on a server—which is potentially a cloud service—and updating the mobile device’s location as shown on floor map running in an app on the mobile device itself.

A thorough understanding of the total delay in updating the client location is essential for the development of effective apps for purposes such as wayfinding. Turn-by-turn navigation, such as that found in outdoor GPS systems, may not be feasible today given the total delay in updating the client location. Wayfinding apps, which show the location of the wireless device on a floor with an arrow providing directions to the indoor destination and periodically update with the new location of the wireless device, may be feasible today. However integration of wireless location technology with additional location technologies such as Bluetooth Low Energy (BLE) may be possible to provide effective turn-by-turn navigation today.

## Apple iOS Version 8 Mobile Devices

As of iOS version 8, Apple changed the way mobile devices send Probe Requests. The basic behavior of Apple iOS 8 devices is that if a device is simply sending Wi-Fi Probe Requests while the mobile device is not being used (for instance the device is in the end user’s pocket or purse), it uses a random fake MAC address. However if the device connects (associates) to the Wi-Fi network, it uses its real MAC address for Probe Requests and for sending packets.

This use of multiple MAC addresses for iOS 8 devices may add to the complexity of tracking such devices when using Probe Requests. However the focus of location-based services in general continues to shift toward engaging the visitor directly within the venue via their mobile device. In general engaging the visitor within the venue has shown to provide the greatest business value. This can be accomplished through the installation of an app on the mobile device and getting the device to connect to the Wi-Fi network within the venue. Enhancements such as the Fastlocate feature which allow all packets to be utilized to determine client location are largely unaffected by the changes to iOS8. The use of an app on the mobile device may also provide the end user the opportunity to “opt-in” to the use of location services within a given venue, helping to further alleviate any potential privacy concerns.

CMX Analytics may simply view an iOS 8 device which uses a random “fake” MAC address to generate Probe Requests as another device for dwell time, crowding, and device count analysis. So there are no issues with the random “fake” MAC address by itself. However when an iOS 8 device uses a random “fake” MAC address to generate Probe Requests and subsequently uses its real MAC address when it associates to the WLAN, it may be viewed as two devices with CMX Analytics. Likewise, the iOS 8 device may use a different random “fake” MAC address when visiting the same venue on multiple days, weeks, etc. In both situations analytics data such as dwell times, crowding, and device counts may be slightly skewed due to double counting. Also CMX Analytics may not be able to view a single person returning with the same iOS 8 device, but using a different random “fake” MAC address as a repeat visitor. Hence repeat visitors data may be slightly skewed. However if the device connects to the Wi-Fi network during both visits, it uses its real MAC address again and CMX Analytics is able to view this as a repeat visitor.

Overall, the following considerations must already be kept in mind regarding CMX Analytics deployments even without iOS 8 devices.



- Not everybody carries a mobile device. If you do not carry a mobile device with Wi-Fi connectivity, you are not counted by CMX Analytics. A given venue will likely have no idea how many people visited the venue without a mobile device. As a result, analytics data may already be skewed slightly based on visitors who do not carry a mobile device with Wi-Fi connectivity.
- Not everybody leaves their Wi-Fi on all the time. The behavior of those who carry mobile devices is often influenced by their cellular data plan. Those with limited voice, but unlimited data plans, may not enable Wi-Fi connectivity since the data transfer through the cellular network has already been paid for. Those with unlimited voice, but limited data plans, may enable Wi-Fi connectivity since the Wi-Fi network may offer a means of accessing the Internet for “free” without having to use their data plan. If you carry a mobile device, but your Wi-Fi is off, you are not counted by CMX Analytics. A given venue will likely have no idea how many people who visited the venue left their Wi-Fi off. As a result, the analytics data may again be skewed slightly.
- Some visitors to a venue may carry multiple mobile devices, each with Wi-Fi connectivity. These visitors are double counted because there is no concept of a single person carrying two or more devices within CMX Analytics. A given venue will likely have no idea how many people who visited were carrying multiple mobile devices. As a result, the analytics data may be skewed slightly.
- If you change out your mobile device—which for the general population is constantly happening since new models are coming out weekly—you end up with a different MAC address which will appear as a different device. So any analytics data regarding repeat visitors may already be skewed slightly due to repeat visitors changing mobile devices over time and the analytics data may be more and more skewed as the timeframe over which you are viewing repeat visitors increases.
- If your battery drops below a certain threshold, your mobile device may not send active Probes at all and you are not counted. A venue will likely have no idea how many visitors had low battery levels and therefore did not send Probe Requests. As a result, the analytics data may be slightly skewed.

The point of this discussion is that analytics data may already be slightly skewed based upon various factors. iOS 8 may be just one additional factor which may skew analytics data slightly more. Just how much further depends upon how many visitors to the venue are carrying iOS 8 mobile devices, whether their Wi-Fi is on, and whether they are connected to Wi-Fi network within the venue. What CMX Analytics provides is the ability to collect information regarding the behavior of a sample of the larger population who visit a venue with the assumption that the behavior of the overall population of visitors is consistent with the sample. Further, CMX Analytics may provide only one of perhaps many data points about customers’ behavior within a venue. Other data points can include direct feedback from staff, sales data, etc. All the data points may be slightly skewed in one way or another. However CMX Analytics may help the venue operator spot trends which can then be acted upon to provide better service to the visitors to the venue.

## Android Mobile Devices

Some mobile devices which run the Android operating system do not generate Probe Requests on 5 GHz Wi-Fi channels which are subject to Dynamic Frequency Selection (DFS). DFS is required by the United States FCC in the U-NII-2 band (channels 52-64, 100-116, and 132-140). Access Points operating on any of these channels may not see Probe Requests from Android devices. Hence location accuracy may be degraded if there are insufficient APs running on non-UNII-2 channels operating to accurately calculate the location of Android devices using only Probe RSSI. Analytics data, which is based upon the location database, may also be slightly skewed because of this.

The network administrator may simply choose to disable U-NII-2 channels on APs via the WLC configuration. However this may reduce the number of available channels in the 5 GHz spectrum from 21 channels to 9 channels when implementing a 20 MHz channel within the U.S. regulatory domain. Similar results may be seen in other regulatory domains. The number of individual channels is further reduced when implementing 40 MHz channels.

In large venues with multiple access points, this may result in the network administrator having to re-use 5 GHz channels. This could result in higher co-channel interference (CCI), reducing the effective throughput of the WLAN. Hence the network administrator must balance location considerations with channel selection and channel width for the optimal solution for the particular venue.

## 2.4 GHz vs. 5 GHz Mobile Devices

Not all mobile devices support both the 2.4 GHz and 5 GHz frequencies. In particular, some older mobile devices support only 2.4 GHz frequencies. Hence when deploying a WLAN for location services within a given venue, it is recommended to deploy the WLAN using both 2.4 GHz and 5 GHz frequency bands. This ensures that devices which operate only in the 2.4 GHz frequency band are seen and location can be calculated for these devices. Since CMX Analytics pulls information from the MSE location database, this helps to optimize analytics data as well.

## FastLocate Deployment Restrictions

The initial version of FastLocate supported has several deployment restrictions.

The FastLocate feature is a global parameter configured on CUWN (AireOS) wireless LAN controllers running software version 8.0. MSE software version 8.0 is also required. The only access points which support FastLocate are the modular Cisco 3600 and 3700 Series APs with Wireless Security Modules (WSMs) installed. The Cisco 3600 and 3700 Series APs also support direct time synchronization via NTP, required for FastLocate to operate. Wireless clients must be connected to the WLAN in order to generate data packets for FastLocate. The APs with WSMs must also handle wireless client traffic.

The MSE is capable of processing RSSI from probe requests as well as data packets. When other (older or non-modular) APs are mixed with Cisco 3600 or 3700 Series APs with WSM modules, the MSE will receive RSSI information from both data packets and probe requests. The MSE supports this deployment, but care needs to be exercised in implementing this type of deployment.

- Best results are achieved when all APs are Cisco 3600 or 3700 Series APs with WSM modules.
- If APs have to be mixed, it helps to achieve some separation between APs capable of processing RSSI from probe requests only and APs capable of processing RSSI from FastLocate (Data RSSI). One way of doing this is using different floors.
- If APs have to be mixed on the same floor try to have zones defined where location refresh requirements are more stringent and other zones that have lower location refresh requirements.
- Do not mix APs in a manner such that every other AP or every 5th AP in the zone is a Cisco 3600 or 3700 Series with a WSM module capable of FastLocate. This could result in unpredictable location results.

FastLocate is currently not supported for FlexConnect deployments with CUWN (AireOS) wireless LAN controllers running software version 8.0.100 and MSE software version 8.0.100.

Given the restrictions discussed above, the following are suggestions as to the possible deployment of FastLocate for location (CAS) services.

- FastLocate is currently not an option for small branches using FlexConnect as shown in Topology #3 in [Figure 4-1](#).
- FastLocate may be a good choice for new standalone sites such as small campuses and large branches which have their own dedicated MSE, WLC, and APs onsite, as shown in Topology #1 in [Figure 4-1](#). In such deployments, it is recommended to deploy Cisco 3700 series APs, since these access points support 802.11ac as well as the WSM module.
- FastLocate may also be a good choice for existing standalone sites such as small campuses and large branches which have their own dedicated MSE, WLC, and APs onsite. In this type of deployment, it may be possible to upgrade all of the access points to Cisco 3600 or 3700 Series APs with WSM modules before implementing location (CAS) services. Again, it is recommended to upgrade to Cisco 3700 series APs, since these access points support 802.11ac as well as the WSM module.
- For standalone sites such as larger campuses with multiple buildings which also have their own dedicated MSE, WLC, and APs onsite, upgrading all of the access points to Cisco 3600 or 3700 Series APs with WSM modules before implementing location (CAS) services may not be feasible due to the amount of time it would take to upgrade all access points in every campus building. Such deployments may also be implementing location (CAS) services through Probe Request RSSI already. As individual floors within the buildings are upgraded to Cisco 3600 or 3700 Series APs with WSM modules, FastLocate may be utilized on a floor-by-floor basis. Again, it is recommended to upgrade to Cisco 3700 series APs since they support 802.11ac as well as the WSM module.





## **PART 3**

### **CMX Use Case Stories**





## CMX Use Case Stories

---

**September 4, 2014**

This chapter introduces several use cases which can be met through the deployment of the CMX services:

- CMX Analytics
- CMX Visitor Connect

Each is designed to highlight the application of CMX services to address a realistic business scenario. The first two use cases involve the use of CMX Location Analytics with a large-sized retail scenario to analyze customer behavior to provide better service. A third use case involves the use of the CMX Visitor Connect service within a small-sized retail scenario to provide customer Wi-Fi access. Finally, a fourth use case involves the use of CMX Presence Analytics within a small-sized retail scenario to analyze customer behavior to provide better service. For each of the use cases, a short Video on Demand (VoD) shows how the use case can be met using the appropriate CMX service.

This initial version of the Cisco CMX CVD focuses on design guidance around proper wireless LAN design for enterprise customers to support location services (CAS), as well as focusing on use case stories around the “detect” aspect of the CMX solution through CMX Location Analytics and CMX Presence Analytics. The design guide also touches upon the “connect” aspect of the CMX solution through a simple CMX Visitor Connect use case story.

Validation work around the Cisco CMX CVD has been done around a series of use case stories. The use case stories are meant to show how CMX services can be used to solve issues or provide benefits in realistic business scenarios within particular vertical markets.




**Note**

A direct link to the VoD for each use case is provided in each of the following sections. If you want to browse the available CMX use case VoDs, see: <http://video.cisco.com/search?q=cmx+cvd+use+case>.

## CMX Location Analytics Use Case Stories

For Location Analytics, two use case stories involving a “big box” retail business scenario were chosen for validation. [Figure 7-1](#) provides an example of the CMX zones applied to the retail use case stories discussed below.

**Figure 7-1 Example of Retail Store Zones for CMX Location Analytics Use Case Stories**



- Use Case Story #1—John is a regional manager for a large big box chain. One of the stores in his region is selling significantly less sporting goods than the rest, while other departments in that store are performing well. He examines the store using CMX Analytics to determine how many people going to that store actually visit the Sporting Goods department using a Conversion Percentage Report. He finds that only about 30% of visitors to that store visit Sporting Goods. For other stores in the region, that percentage averages 45%. Since the Sporting Goods department is near the rear of the store, he directs the store manager to add signage at the entrance promoting sporting goods. A week and a half later, John returns to CMX Analytics to track whether the changes that were made were effective. He notices that the average percentage goes up to around 40% for that store, thus improving the Sporting Goods department conversion rate. John can also look at conversion rates from each of the other departments in relation to Sporting Goods to help further fine tune changes that increase the conversion rate and subsequently sales.
- Use Case Story #2—John digs deeper into CMX Analytics and examines the newest store in his region, which opened two months ago. He discovers an issue in the Appliances & Kitchenware department. Using the Analytics interface, he sees that dwell time is much lower than it should be, showing an average of less than eight minutes, where the regional average is closer to 15 minutes. Even without having the initial sales data from the store, John knows improvements can be made and works with the store manager to initiate an immediate training session for the staff in that department, focusing on customer engagement. They both then track how much impact the training had and see that average dwell time has increased to 18 minutes, surpassing the regional average. The training is then implemented region wide due to its effectiveness.

The following short VoD shows how the Conversion Percentage Report within Cisco CMX Analytics can be used to provide the information that the regional manager needs to meet the first use case and how the CMX Analytics Dashboard can be used to provide the information he needs to meet the second use case.

<http://video.cisco.com/detail/video/3753883722001/cisco-cmx-cvd-use-case-stories-1-and-2?autoStart=true>



## CMX Visitor Connect Use Case Story

For CMX Visitor Connect, a use case story involving a specialty retail business scenario was chosen for validation.

Suyog is the owner/manager of a small business, “Suyog’s Chai House”. Suyog wishes to provide guest WiFi Internet access for his customers using a simple, easy, and customizable guest captive portal to interface with the guest for onboarding. He wishes to provide the ability for his guests to login via Facebook (priority), Linked-In, and Google+ (stretch goals), collect the name and an email address of his guests for targeted email promotions, and display terms and conditions for using the guest WiFi Internet access. However he needs to also provide WiFi guest access for those customers who wish to join anonymously. Suyog wishes to push a promotional advertisement and coupon for “Suyog’s Special Chai Blend” tea as customers log in via social media and redirect them to his website when they first connect. Optionally, as an incentive for customers to login via social media, Suyog will provide unlimited daily access to customers. For those customers who login anonymously, Suyog will enforce a quota limit. Suyog wants to make sure that customers who temporarily disassociate and re-associate to the network do not have to log in again via social media.

The following short VoD shows how the CMX Visitor Connect service can be used by the owner/manager of a small business to provide the guest WiFi access he needs to meet this use case.

[VoD to be provided at future date.]

## CMX Presence Analytics Use Case Story

For CMX Presence Analytics, a use case story involving another specialty retail business scenario was chosen for validation.

Conrad is the owner/manager of a small business, “Conrad's Cups and Cakes” specializing in coffee and cupcakes. Conrad wishes to provide guest Wi-Fi Internet access for his customers. In exchange for free Wi-Fi service, Conrad is interested in collecting analytics information for his store as well. Specifically, he is interested in knowing the percentage of customers passing by who actually stop in his stores, presumably to buy something. Likewise, he would like information regarding the average number of customers within his stores during various times of the day to provide sufficient staff for a superior level of service to his customers. Finally, he is interested in length of time customers spend in his stores.

The following short Video on Demand (VoD) shows how the CMX Analytics Dashboard can be used for a presence analytics site to provide the owner/manager of the small business the information he needs to meet this use case.

<http://video.cisco.com/detail/video/376978000001/cisco-cmx-use-case-story-4?autoStart=true>





## **PART 4**

# **CMX Radio Frequency and Location Based Design**





# Summary of CMX Radio Frequency and Location Based Design

---

September 4, 2014

This part of the CVD includes the following chapters:

- [Chapter 9, “Radio Operating Frequencies and Data Rates”](#)—Discusses radio operating frequencies that are used for WLAN networks. Also briefly discussed are 802.11 a/b/n/ac modulation techniques and the role of TPC and DCA in a radio frequency (RF) network.
- [Chapter 10, “Radio Frequency Fundamentals”](#)—Explains various RF concepts like spectrum bands, power level, signal strength, RSSI, etc. and provides a simple example using these concepts to illustrate how RF impacts a client and Access Point perceived signal strength level. These components include Cisco wireless LAN controllers (WLCs), Cisco Prime Infrastructure (PI), and the Cisco Mobility Services Engine (MSE). In addition, the configuration of CMX services, specifically CMX Analytics and CMX Visitor Connect, are discussed.
- [Chapter 11, “Antenna Fundamentals”](#)—Discusses antennas, which are a fundamental part of any WLAN deployment since selecting the right type of antenna for deployment greatly enhances not just coverage, but also location readiness.
- [Chapter 12, “802.11 Fundamentals”](#)—Discusses 802.11 Fundamentals, namely the role of beacons, probe requests, and probe responses.
- [Chapter 13, “Location Fundamentals”](#)—This part of the CVD discusses location fundamentals, including the definition of a location ready point, location currency, location accuracy, and location latency. We also discuss two methods of getting location from a client, i.e., the Probe RSSI method and the FastLocate method.
- [Chapter 14, “Pre-Deployment Radio Frequency Site Survey”](#)—A good Cisco WLAN deployment is dependent on a good RF design, which includes doing a thorough site survey of the location, determining the best location for access points, making the right channel plans, planning for AP capacity, and lastly performing a regular post deployment RF site survey. The chapter discusses pre-deployment site survey topics.
- [Chapter 15, “Access Point Placement and Separation”](#)—Discusses AP placement and AP capacity planning, including core concepts regarding the distance between APs in a network and its impact on location, data, and voice.
- [Chapter 16, “Predictive Radio Frequency Planning”](#)—Discusses predictive RF planning that should be undertaken after a pre-deployment RF site survey is completed and two tools to perform RF planning, the Cisco Prime Infrastructure RF Planner tool and the Ekahau Site Survey tool.

- [Chapter 17, “Multi-Floor Deployments”](#)—Discusses challenges in deployments that involve multiple floors. Recommendations on what to keep in mind while designing for RF network are also discussed.
- [Chapter 18, “Capacity Planning and High Density”](#)—Discusses planning a network while keeping capacity and application requirements in mind.
- [Chapter 19, “Location Voice and Data Co-Existence”](#)—Discusses the pertinent characteristics of voice and data designs only as they relate to co-existence with the location tracking capabilities of the Cisco UWN.
- [Chapter 20, “Post-Deployment Radio Frequency Tuning”](#)—Discusses post deployment RF tuning that should be done regularly on the deployment and includes using RRM for channel planning, CleanAir to mitigate RF interference, and a regular post-site survey assessment to ensure that optimum RF health is maintained.
- [Chapter 21, “Best Practices Checklist”](#)—Discusses the best practices check list while deploying a CMX solution.



# Radio Operating Frequencies and Data Rates

September 4, 2014

This part of the CVD discusses radio operating frequencies that are used for WLAN networks. Also briefly discussed are 802.11 a/b/n/ac modulation techniques and the role of TPC and DCA in a radio frequency (RF) network.

## Radio Frequency Bands

In the United States there are three radio frequency bands allocated for unlicensed industrial, scientific, and medical (ISM) usage, as shown in [Table 9-1](#).

**Table 9-1** Radio Frequency Bands

RF Band	Uses
900 MHz (902 to 928 MHz)	Phones, medical equipment, RFID, etc.
2.4 GHz (2.4 to 2.4835 GHz)	IEEE 802.11b/g/n, Bluetooth, microwave ovens, analog video cameras, older DECT standard-based phones, etc.
5 GHz (5.15 to 5.35 and 5.725 to 5.825 GHz)	IEEE 802.11a/n and 802.11ac, analog video cameras, radars, older DECT standard-based phones, etc.

For our purpose we are interested in the IEEE 802.11b/n, 802.11a/n, and 802.11ac for Wi-Fi communication. Each range has different characteristics. The lower frequencies exhibit better range, but with limited bandwidth and thus lower data rates. The higher frequencies exhibit less range and are subject to greater attenuation from solid objects.

## Regulatory Domains

Devices that operate in unlicensed bands do not require any formal licensing process, but when operating in these bands, the user is obligated to follow the government regulations for that region. The regulatory domains in different parts of the world monitor these bands according to different criteria and the WLAN devices used in these domains must comply with the specifications of the relevant governing regulatory domain. Although the regulatory requirements do not affect the interoperability of IEEE 802.11b/g/n and 802.11a/n/ac-compliant products, the regulatory agencies do set certain criteria in the standard, such as specifying the emission requirements for WLAN to minimize the amount of interference a radio can

generate or receive from another radio in the same proximity. It is the responsibility of the vendor to obtain certification for a product from the relevant regulatory body. Table 9-2 summarizes the current regulatory domains for Wi-Fi products, which include the main regulatory domains FCC, ETSI, and MKK.

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

**Table 9-2 Regulatory Domains**

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

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



**Note**

See

[http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1300-series/product\\_data\\_sheet0900aecd80537b6a.html](http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1300-series/product_data_sheet0900aecd80537b6a.html) for compliance information and also check with your local regulatory authority to find out what is permitted within your country. The information provided in Table 9-1 and Table 9-2 should be used as general guidelines.

## Operating Frequencies

The 2.4 GHz band regulations have been relatively constant, given the length of time they have been operating. The FCC allows for 11 channels, ETSI allows for up to 13 channels, and Japan allows up to 14 channels, but requires a special license to operate in channel 14.

For 5 GHz, countries are moving to open the frequency range 5.250-5.350 GHz (UNII-2) and the frequency range 5.470 to 5.780 GHz for additional 802.11a channels. These various frequencies are covered in more detail in the specific 802.11 sections in this chapter.

## 802.11 Modulation Techniques

The IEEE 802.11 standard makes provisions for the use of several different modulation techniques to encode the transmitted data onto the RF signal. These modulation techniques are used to enhance the probability of the receiver correctly receiving the data and thus reducing the need for retransmissions. The techniques vary in their complexities and robustness to RF signal propagation impairments.



## Direct-Sequence Spread Spectrum

The direct-sequence spread spectrum (DSSS) approach involves encoding redundant information into the RF signal. Every data bit is expanded to a string of chips called a chipping sequence or Barker sequence. The chipping rate, as mandated by the U.S. FCC, is 10 chips at the 1- and 2-Mbps rates and 8 chips at the 11-Mbps rate. So, 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 channel.

## Frequency-Hopping Spread Spectrum

Frequency-hopping spread spectrum (FHSS) uses a radio that moves or hops from one frequency to another at predetermined times and channels. The regulations require that the maximum time spent on any one channel is 400 milliseconds. For the 1- and 2-Mb FHSS systems, the hopping pattern must include 75 different channels and must use every channel before reusing any one. For wide-band frequency hopping (WBFH) systems, which permit up to 10-Mb data rates, the rules require the use of at least 15 channels and they cannot overlap. With only 83 MHz of spectrum, WBFH limits the systems to 15 channels, thus causing scalability issues. In every case, for the same transmitter power and antennas, a DSSS system has greater range, scalability, and throughput than an FHSS system. For this reason, Cisco has chosen to support only direct-sequence systems in the spread spectrum products.

## Orthogonal Frequency Division Multiplexing

The orthogonal frequency division multiplexing (OFDM) used in 802.11a/n/ac and 802.11g data transmissions offers greater performance than the older direct-sequence systems. In the OFDM system, each tone is orthogonal to the adjacent tones and therefore does not require the frequency guard band needed for direct sequence. This guard band lowers the bandwidth efficiency and wastes up to 50 percent of the available bandwidth. Because OFDM is composed of many narrow-band tones, narrow-band interference degrades only a small portion of the signal, with little or no effect on the remainder of the frequency components.

**Note**

802.11 beacon frames are sent at the highest power at lowest possible data rate. For example, in a 2.4GHz 802.11n deployment, if all data rates are enabled, then beacons are sent at highest power at 1 Mbps data rate. At 1 Mbps, DSSS is used instead of OFDM. In a good WLAN deployment it is advisable to disable data rates that are not used. More deployment guidelines are discussed in later sections of this design guide.

## 2.4 GHz Operating Frequencies and Data Rates

Ratified in September, 1999, the 802.11b standard operates in the 2.4 GHz spectrum and supports data rates of 1, 2, 5.5, and 11 Mbps. 802.11b enjoyed broad user acceptance and vendor support for a number of years.

The 802.11g standard, which was ratified in June, 2003, operates in the same spectrum as 802.11b and is backward-compatible with the 802.11b standard. 802.11g supports the additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps.

The 802.11n standard, which was ratified in 2009, also operates in the same spectrum as 802.11b/g and is backward-compatible with the 802.11b standard. 802.11n made a fundamental shift in WLAN technology by introducing many new improvements over older 802.11b/g technology to provide much higher speeds. These new improvements include channel bonding to provide 40 MHz channels, MIMO (Multiple In Multiple Out) Antenna system, MRC (Maximum Ratio Combining) and Shorter Guard Intervals. The data rates in the 2.4 GHz spectrum can reach up to 72 Mbps with a single channel (20 MHz), single antenna mode. With a 20 MHz channel and multiple antennas, 802.11n data rates in the 2.4 GHz can reach up to 300 Mbps. By combining two 20 MHz channels, throughputs up to 600MHz can be achieved, however channel bonding is not a popular choice in a 2.4 GHz deployment due to various factors like RF interference and lack of overlapping channels.

**Note**

While 802.11n technology is supported in 2.4 GHz spectrum, Cisco Access Points cannot be configured to use 40 MHz channels in the 2.4 GHz spectrum by combining channels.

**Note**

The operating frequencies, modulation schemes and bands supported tables are discussed in [Appendix C, “802.11 Data Rates.”](#)

## 5 GHz Operating Frequencies and Data Rates

Operating in the unlicensed portion of the 5 GHz radio band, 802.11a is reasonably 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 standard operates in a different frequency range, it is not compatible with existing 802.11b or 802.11g-compliant wireless devices, but it does mean that 2.4-GHz and 5-GHz equipment can operate in the same physical environment without interference.

Choosing between these two technologies (802.11b/g/n and 802.11a/n/ac) does not involve a one-for-one trade-off. They are complementary technologies and will continue to coexist in future enterprise environments. Those responsible for implementing these technologies must be able to make an educated choice to properly plan 2.4 GHz and 5 GHz networks. Proper WLAN capacity planning should be undertaken to ensure that devices both in 2.4 GHz and 5 GHz get adequate bandwidth for their use.

WLAN capacity planning is discussed in more detail in [Chapter 15, “Access Point Placement and Separation.”](#)

The 5 GHz band in which 802.11a/n/ac operates is divided into several different sections. Each of the Unlicensed National Information Infrastructure (UNII) bands presented in [Table 9-3](#) was originally intended for different uses, but all can currently be used by indoor 802.11a/n/ac 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 using 20 MHz channel width.

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

[Table 9-3](#) summarizes the channels available and restrictions on them.

**Table 9-3 Channels Available**

UNII band	Channels Available	Restriction
UNII-1, 5.150 to 5.250 GHz	36, 40, 44, 48	Not Applicable
UNII-2, 5.250 to 5.350 GHz	52, 56, 60, 64	Requires Dynamic Frequency Selection (DFS) to avoid Radars and Transmitter Power Control (TPC)
UNII-2 5.470 to 5.725 GHz	100, 104, 108, 112, 116, 120, 124, 128, 132, 136, 140	Requires Dynamic Frequency Selection (DFS) to avoid Radars and Transmitter Power Control (TPC)
UNII-3 5.725 to 5.825 GHz	149, 153, 157, 161, 165	Don't require DFS or TPC

**Note**

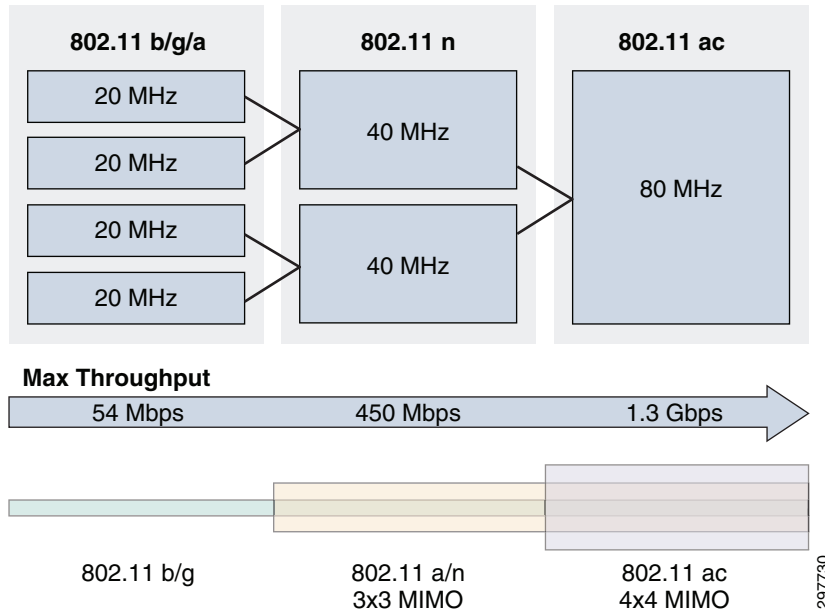
Not all channels in a given range can be used in all of the regulatory domains. Refer to [“802.11 Data Rates”](#) for information on supported channel and regulatory domain.

802.11a uses 20 Mhz channels. 802.11n can use 20 or 40 MHz channels. 802.11ac is defined up to 160 MHz channels. However current implementations (802.11ac Wave 1 products) can use 20, 40, or 80 MHz channels. The way channels are used is through channel bonding, i.e., taking adjacent channels and bonding them together to create a bigger spectrum (channel) that facilitates higher data rate. The 802.11n standard, which was ratified in 2009, also operates in the same spectrum as 802.11a and is backward compatible with the 802.11a standard. 802.11n made a fundamental shift in WLAN technology by introducing many new improvements over older 802.11a technology to provide much higher speeds. These new improvements include channel bonding to provide 40 MHz channels, MIMO (Multiple In Multiple Out) Antenna system, MRC (Maximum Ratio Combining), and Shorter Guard Intervals. With a 20 MHz channel and multiple antennas, 802.11n data rates in the 5 GHz can reach up to 300 Mbps. By combining two 20 MHz channels into a 40MHz channels, throughputs up to 600MHz can be achieved. Similarly by combing even more channels into an 80Mhz channels for 802.11ac, throughputs up to 1.3Gb can be achieved.

**Note**

There are not many commercial implementations of APs which go beyond 3x3 MIMO with three spatial streams however. Therefore you are unlikely to currently find an AP with higher than 450 Mbps data rates for 802.11n.

Figure 9-1 Comparison of Different Protocols

**Note**

A good discussion on 802.11n data rates can be found at:

[http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1250-series/design\\_guide\\_c07-693245.html#wp9001157](http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1250-series/design_guide_c07-693245.html#wp9001157).

A note on the DFS and TPC: Dynamic Frequency Selection 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, listening 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. The transmitting device continuously monitors the environment for the presence of radar, both prior to and during operation. 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 collocated.

TPC allows the AP to negotiate power levels with a WLAN client during that association process. The AP can inform that WLAN client of the range of allowable transmit power to be used with that AP and may 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.

**Note**

Both DFS and TPC are discussed in greater detail in [Radio Resource Management](#) in [Chapter 20](#), “Post-Deployment Radio Frequency Tuning.”

## 802.11ac

802.11ac is a faster and more scalable version of 802.11n. 802.11ac couples the freedom of wireless with the capabilities of Gigabit Ethernet. 802.11ac achieves its raw speed increase by pushing on three different dimensions:

- More channel bonding, increased from the maximum of 40 MHz in 802.11n and now up to 80 or even 160 MHz (for 117% or 333% speed-ups, respectively).
- Denser modulation, now using 256 quadrature amplitude modulation (QAM), up from 802.11n's 64QAM (for a 33% speed burst at shorter, yet still usable, ranges).
- More multiple inputs, multiple outputs (MIMO). Whereas 802.11n stopped at four spatial streams, 802.11ac goes all the way to eight (for another 100% speed-up).

**Note**

---

Current Wave 1 of 802.11ac supports only up to 4x4 spatial streams.

---

The design constraints and economics that kept 802.11n products at one, two, or three spatial streams have not changed much for 802.11ac, so we can expect the same kind of product availability with first-wave 802.11ac products built around 80 MHz and delivering up to 433 Mbps (low end), 867 Mbps (mid-tier), or 1300 Mbps (high end) at the physical layer. Second-generation products promises still more channel bonding and spatial streams, with plausible product configurations operating at up to 3.47 Gbps.

802.11ac is a 5 GHz-only technology, so dual-band APs and clients will continue to use 802.11n at 2.4 GHz. However 802.11ac clients operate in the less crowded 5 GHz band.

**Note**

---

Current Wave 1 of 802.11ac implementations support up to 80 MHz channels only. Wave 2 of 802.11ac implementation will support up to 160 Mhz channels.

---

**Note**

---

Because of increased width, there are now far fewer channels available for 802.11ac in 5 GHz. To get maximum throughput both client and AP should support 802.11ac. In general you also need good SNR ratio across all four sub 20 Mhz channels to achieve highest throughput.

---

**Note**

---

For a more detailed discussion around 802.11ac, see:  
[http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-3600-series/white\\_paper\\_c11-713103.html](http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-3600-series/white_paper_c11-713103.html).

---





# Radio Frequency Fundamentals

---

September 4, 2014

This part of the CVD discusses Radio Frequency (RF) fundamentals that are necessary to understand before deploying a Wireless LAN network that is location and CMX ready. The chapter explains various RF concepts like spectrum bands, power level, signal strength, RSSI, etc. and provides a simple example using these concepts to illustrate how RF impacts a client and Access Point perceived signal strength level. These components include Cisco wireless LAN controllers (WLCs), Cisco Prime Infrastructure (PI), and the Cisco Mobility Services Engine (MSE). In addition, the configuration of CMX services, specifically CMX Analytics and CMX Visitor Connect, are discussed.

A good WLAN deployment either in the 2.4 GHz spectrum or the 5 GHz spectrum is contingent on Radio Frequency Planning. As we noted, because WLAN technology operates in the unlicensed bands provided by the FCC, many other wireless technologies such as Bluetooth also use the same spectrum. It is important for WLAN deployments to consider RF characteristics, regulatory domains, maximum transmittable power by APs, and RF interference from other bands. This chapter explains the basic terminologies used in radio frequencies and how they tie into the WLAN deployment. General guidelines on designing a WLAN with RF in mind are covered in [Chapter 14, “Pre-Deployment Radio Frequency Site Survey.”](#)

## Power Level

The dB measures the power of a signal as a function of its ratio to another standardized value. The abbreviation dB is often combined with other abbreviations to represent the values that are compared. Here are two examples:

- dBm—The dB value is compared to 1 mW.
- dBw—The dB value is compared to 1 W.

You can calculate the power in dBs with the formula:

$$\text{Power (in dB)} = 10 * \log_{10} (\text{Signal/Reference})$$

The terms in the formula:

- log<sub>10</sub> is logarithm base 10.
- Signal is the power of the signal (for example, 50 mW).
- Reference is the reference power (for example, 1 mW).

For example, if you want to calculate the power in dB of 50 mW, apply the formula to obtain:

$$\text{Power (in dB)} = 10 * \log_{10} (50/1) = 10 * \log_{10} (50) = 10 * 1.7 = 17 \text{ dBm}$$

Because decibels are ratios that compare two power levels, you can use simple math to manipulate the ratios for the design and assembly of networks. For example, you can apply this basic rule to calculate logarithms of large numbers:

$$\log_{10}(A*B) = \log_{10}(A) + \log_{10}(B)$$

If you use the formula above, you can calculate the power of 50 mW in dBs in this way:

$$\text{Power (in dB)} = 10 * \log_{10}(50) = 10 * \log_{10}(5 * 10) = (10 * \log_{10}(5)) + (10 * \log_{10}(10)) = 7 + 10 = 17 \text{ dBm}$$

Table 10-1 lists some commonly used general rules.

**Table 10-1** *db Values and tx Power*

An Increase of:	A Decrease of:	Produces:
3 dB		Double transmit power
	3 dB	Half transmit power
10 dB		10 times the transmit power
	10 dB	Divides transmit power by 10
30 dB		1000 times the transmit power
	30 dB	Decreases transmit power 1000 times

Table 10-2 provides approximate dBm to mW values.

**Table 10-2** *dBm to mW Values*

dBm	mW
0	1
1	1.25
2	1.56
3	2
4	2.5
5	3.12
6	4
7	5
8	6.25
9	8
10	10
11	12.5
12	16
13	20
14	25
15	32
16	40



**Table 10-2** *dBm to mW Values*

dBm	mW
17	50
18	64
19	80
20	100
21	128
22	160
23	200
24	256
25	320
26	400
27	512
28	640
29	800
30	1000 or 1 W

Here is an example:

If 0 dB = 1 mW, then 14 dB = 25 mW.

If 0 dB = 1 mW, then 10 dB = 10 mW, and 20 dB = 100 mW.

Subtract 3 dB from 100 mW to drop the power by half (17 dB = 50 mW). Then, subtract 3 dB again to drop the power by 50 percent again (14 dB = 25 mW).



**Note**

You can find all values with a little addition or subtraction if you use the basic rules of algorithms.

## Effective Isotropic Radiated Power

The radiated (transmitted) power is rated in either dBm or W. Power that comes off an antenna is measured as effective isotropic radiated power (EIRP). EIRP is the value that regulatory agencies, such as the FCC or European Telecommunications Standards Institute (ETSI), use to determine and measure power limits in applications such as 2.4-GHz or 5-GHz wireless equipment. To calculate EIRP, add the transmitter power (in dBm) to the antenna gain (in dBi) and subtract any cable losses (in dB).

Table 10-3 shows an example.

**Table 10-3** *Tx Power Relationship*

Part	Cisco Part Number	Power
A Cisco Access Point	AIR-CAP3702I	20 dBm
That uses a 50 foot antenna cable	AIR-CAB050LL-R	3.35 dB loss

**Table 10-3 Tx Power Relationship**

Part	Cisco Part Number	Power
And a solid dish antenna	AIR-ANT3338	21 dBi gain
Has an EIRP of		37.65 dBm (20 - 3.35 + 21)

## Path Loss

The distance that a signal can be transmitted depends on several factors. The primary hardware factors that are involved are:

- Transmitter power (in dBm)
- Cable losses between the transmitter and its antenna (in dB)
- Antenna gain of the transmitter (in dBi)
- Signal Attenuation—This refers to how far apart the antennas are and if there are obstacles between them.
- Receiving antenna gain—Cable losses between the receiver and its antenna
- Receiver sensitivity

## Receive Signal Strength Indicator—RSSI

Receiver sensitivity is defined as the minimum signal power level with an acceptable Bit Error Rate (in dBm or mW) that is necessary for the receiver to accurately decode a given signal. This is usually expressed in a negative number depending on the data rate. For example an Access Point may require an RSSI of at least negative -91 dBm at 1 MB and an even higher strength RF power -79 dBm to decode 54 MB.

Because 0 dBm is compared to 1 mW, 0 dBm is a relative point, much like 0 degrees is in temperature measurement. [Table 10-4](#) shows example values of receiver sensitivity.

**Table 10-4 dBm to mW Receiver Sensitivity**

dBm	mW
10	10
3	2
0	1
-3	0.5
-10	0.1
-20	0.01
-30	0.001
-40	0.0001
-50	0.00001
-60	0.000001

**Table 10-4** *dBm to mW Receiver Sensitivity*

dBm	mW
-70	0.0000001
-84	0.000000004

## Signal to Noise Ratio—SNR Ratio

Noise is any signal that interferes with your signal. Noise can be due to other wireless devices such as cordless phones, microwave devices, etc. This value is measured in decibels from 0 (zero) to -120 (minus 120). Noise level is the amount of interference in your wireless signal, so lower is generally good for WLAN deployments. Typical environments range between -90dBm and -98dBm with little ambient noise. This value may be even higher if there is a lot of RF interference coming in from other non-802.11 devices on the same spectrum

Signal to Noise Ratio or SNR is defined as the ratio of the transmitted power from the AP to the ambient (noise floor) energy present. To calculate the SNR value, we add the Signal Value to the Noise Value to get the SNR ratio. A positive value of the SNR ratio is always better.

For example, say your Signal value is -55dBm and your Noise value is -95dBm.

The difference of Signal (-55dBm) + Noise (-95dBm) = 40db—This means you have an SNR of 40.

Note that in the above equation you are not merely adding two numbers, but are interested in the “difference” between the Signal and Noise values, which is usually a positive number. The lower the number, the lower the difference between noise and transmitted power, which in turn means lower quality of signal. The higher the difference between Signal and Noise means that the transmitted signal is of much higher power than the noise floor, thereby making it easier for a WLAN client to decode the signal.

The general rule of thumb for a good RF deployment of WLAN is that any SNR should be above 20-25.

## Signal Attenuation

Signal attenuation or signal loss occurs even as the signal passes through air. The loss of signal strength is more pronounced as the signal passes through different objects. A transmit power of 20 mW is equivalent to 13 dBm. Therefore if the transmitted power at the entry point of a plasterboard wall is at 13 dBm, the signal strength will be reduced to 10 dBm when exiting that wall. Some common examples are shown in [Table 10-5](#).

**Table 10-5** *Material and Signal Attenuation*

Material/Object	Signal Attenuation <sup>1</sup>
Plasterboard wall	3 dB
Glass wall with metal frame	6 dB
Cinder block wall	4 dB
Office window	3 dB
Metal door	6 dB

**Table 10-5** Material and Signal Attenuation

Material/Object	Signal Attenuation <sup>1</sup>
Metal door in brick wall	12 dB
Human Body	3 dB

1. Values are approximate.

## Example Use Case

Here is an example to tie together this information to come up with a very simple RF plan calculator for a single AP and a single client.

- Access Point Power = 20 dBm
- 50 foot antenna cable = - 3.35 dB Loss
- Signal attenuation due to glass wall with metal frame = -6 dB
- External Access Point Antenna = + 5.5 dBi gain
- RSSI at WLAN Client = -75 dBm at 100ft from the AP
- Noise level detected by WLAN Client = -85 dBm at 100ft from the AP

Based on the above, we can calculate the following information.

- EIRP of the AP at source =  $20 - 3.35 + 5.5 = 22.15$  dBm
- Transmit power as signal passes through glass wall =  $22.15 - 6 = 16.15$  dBm
- SNR at Client =  $-75 + -85 = 10$  dBm (difference between Signal and Noise)

We can see that an SNR of just 10 means a low quality signal connection between the AP and client. To correct this, there are several options:

- Move the client closer to AP, thereby increasing the RSSI at the client, which in turn gives a better SNR ratio.
- Increase the power transmitted from the AP, which increases the RSSI at the WLAN client.
- Increase the power transmitted with help of a higher gain antenna, which increases RSSI at WLAN client.
- Remove sources of interferences in the WLAN area to reduce the Noise level, thereby increasing the SNR at WLAN Client.
- Increase AP density by putting in an AP nearer to client, which increases RSSI at client and improve SNR ratio.

As we see, there is no one size fits all solution. Similarly, a good RF deployment involves a proper RF site survey to plan for good coverage, link estimation, and capacity. This includes planning for interference, materials and physical structure of the space, antennas, and power levels. Further sections provide details about a location-ready design.



# Antenna Fundamentals

---

September 4, 2014

This part of the CVD discusses antennas, which are a fundamental part of any WLAN deployment since selecting the right type of antenna for deployment greatly enhances not just coverage, but also location readiness.

An antenna gives the wireless system three fundamental properties—gain, direction, and polarization. Gain is a measure of increase in power. Direction is the shape of the transmission pattern. A good analogy for an antenna is the reflector in a flashlight. The reflector concentrates and intensifies the light beam in a particular direction similar to what a parabolic dish antenna would do to a RF source in a radio system.

## Antenna Gain

Antenna gain is measured in decibels, which is a ratio between two values. The gain of a specific antenna is compared to the gain of an isotropic antenna. An isotropic antenna is a theoretical antenna with a uniform three-dimensional radiation pattern (similar to a light bulb with no reflector). dBi is used to compare the power level of a given antenna to the theoretical isotropic antenna. The U.S. FCC uses dBi in its calculations. An isotropic antenna is said to have a power rating of 0 dB, meaning that it has zero gain/loss when compared to itself.

Unlike isotropic antennas, dipole antennas are real antennas. Dipole antennas have a different radiation pattern compared to isotropic antennas. The dipole radiation pattern is 360 degrees in the horizontal plane and 75 degrees in the vertical plane (assuming the dipole antenna is standing vertically) and resembles a donut in shape. Because the beam is “slightly” concentrated, dipole antennas have a gain over isotropic antennas of 2.14 dB in the horizontal plane. Dipole antennas are said to have a gain of 2.14 dBi (in comparison to an isotropic antenna).

Some antennas are rated in comparison to dipole antennas, which is denoted by the suffix dBd. Hence dipole antennas have a gain of 0 dBd (= 2.14 dBi).



### Note

---

Majority of documentation refers to dipole antennas as having a gain of 2.2 dBi. The actual figure is 2.14 dBi, but is often rounded up.

---

You can also use the dB abbreviation to describe the power level rating of antennas:

- dBi—For use with isotropic antennas.
- dBd—With reference to dipole antennas.

The power rating difference between dBd and dBi is approximately 2.2—that is, 0 dBd = 2.2 dBi. Therefore an antenna that is rated at 3 dBd is rated by the FCC (and Cisco) as 5.2 dBi.


## Antenna Types

Cisco offers several different styles of antennas for use with access points and bridges in both 2.4-GHz and 5-GHz products. Every antenna offered for sale has been FCC approved. Each type of antenna offers different coverage capabilities. As the gain of an antenna increases, there is some tradeoff to its coverage area. Usually high-gain antennas offer longer coverage distances, but only in a certain direction. The radiation patterns below help to show the coverage areas of the styles of antennas that Cisco offers: omnidirectional, Yagi, and patch antennas.

## Omnidirectional Antenna

An omnidirectional antenna ([Figure 11-1](#)) is designed to provide a 360 degree radiation pattern. This type of antenna is used when coverage in all directions from the antenna is required. The standard 2.14-dBi “Rubber Duck” is one style of omnidirectional antenna.


**Figure 11-1**      *Omnidirectional Antennas*




## Directional Antennas

Directional antennas come in many different styles and shapes. An antenna does not offer any added power to the signal; it simply redirects the energy it receives from the transmitter. By redirecting this energy, it has the effect of providing more energy in one direction and less energy in all other directions. As the gain of a directional antenna increases, the angle of radiation usually decreases, providing a greater coverage distance, but with a reduced coverage angle. Directional antennas include patch antennas ([Figure 11-2](#)), Yagi antennas ([Figure 11-3](#)), and parabolic dishes. Parabolic dishes have a very narrow RF energy path and the installer must be accurate in aiming these types of antennas these at each other.

**Figure 11-2 Directional Antenna**



**Figure 11-3 Yagi Antenna**




# Multipath Distortion

Multipath interference occurs when an RF signal has more than one path between a receiver and a transmitter. This occurs in sites that have a large amount of metallic or other RF reflective surfaces.

Just as light and sound bounce off of objects, so does RF. This means there can be more than one path that RF takes when going from a transmit (TX) and receive (RX) antenna. These multiple signals combine in the RX antenna and receiver to cause distortion of the signal.

Multipath interference can cause the RF energy of an antenna to be very high, but the data would be unrecoverable. Changing the type of antenna and location of the antenna can eliminate multipath distortion (Figure 4).

**Figure 11-4 Multipath Distortion**



You can relate multipath distortion to a common occurrence in your car. As you pull up to a stop, you may notice static on the radio. But as you move forward a few inches or feet, the station starts to come in more clearly. By rolling forward, you move the antenna slightly, out of the point where the multiple signals converge.

## Diversity Antenna Systems and Multipath Distortion

A diversity antenna system can be compared to a switch that selects one antenna or another, but never both at the same time. The radio in receive mode continually switches between antennas listening for a valid radio packet. After the beginning synchronization if a valid packet is heard, the radio evaluates the synchronization signal of the packet on one antenna and then switches to the other antenna and evaluates. Then the radio selects the best antenna and uses only that antenna for the remaining portion of that packet.

On Transmit, the radio selects the same antenna it used the last time it communicated with that given radio. If a packet fails, it switches to the other antenna and retries the packet.

Diversity antenna systems are used to overcome a phenomenon known as multipath distortion or multipath interference. A diversity antenna system uses two identical antennas located a small distance apart to provide coverage to the same physical area.

One caution with diversity antenna systems is that they are not designed for using two antennas covering two different coverage cells. The problem in using it this way is that if antenna number 1 is communicating to device number 1 while device number 2 (which is in the antenna number 2 cell) tries to communicate, antenna number 2 is not connected (due to the position of the switch) and the communication fails. Diversity antennas should cover the same area from only a slightly different location.

With the introduction of the latest direct-spread spectrum physical layer chips and the use of diversity antenna systems, direct-spread spectrum systems have equaled or surpassed frequency-hopping systems in handling multipath interference. While the introduction of Wide Band Frequency Hopping does increase the bandwidth of frequency-hopping systems, it drastically affects the ability to handle multipath issues, further reducing its range compared to present direct-spread systems in sites that are highly RF reflective.

## Antenna Orientation and Access Point Placement

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


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

Antenna Orientation ([Figure 11-5](#)) illustrates how the configuration of the antenna's azimuth orientation within Cisco Prime Infrastructure is mapped to the actual physical orientation of the antenna. The blue triangle in the azimuth compass rose shown at the right of the figure indicates how the actual antenna should be physically positioned during deployment (notice that each of the antenna graphics contains a blue arrow as well). For omni-directional antennas, use unique identifying factors that are associated with the antenna (such as the right angled flexible antenna connector shown at the bottom of the 2.2dBi



black whip antenna in [Figure 11-5](#)) to assist in proper positioning. For directional antennas, use unique physical characteristics of the antenna such as the exit location of the cable (for example, cable exiting up or cable exiting down) or other unique marks and construction characteristics.

**Figure 11-5** Antenna Orientation



In software Release 7.0 and above of the location aware Cisco UWN, the ability to specify installed access point and antenna characteristics has been enhanced. CUWN provides the ability to account for antenna installations at varying heights for individual APs. This capability allows the Mobility Services Engine (MSE) to incorporate varied access point antenna heights in its lateration calculations, an approach that more realistically approximates installations, especially in non-carpeted office type environments. Note that the individual height specified for an access point antennas cannot exceed the height of the floor.

## Defining Individual Access Point Heights

**Figure 11-6** Access Point Heights

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


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

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

The antenna propagation characteristics of the AP3700 access point are optimal along its azimuth plane when ceiling mounted. For optimal location performance when using the AP3700, it is preferred that the access point be ceiling mounted rather than wall mounted.

In some cases, it is desirable to separate access points from antennas using a short length (less than 10 feet) of low-loss antenna cable. Reasons for this might include avoidance of obstacles or simply the desire to position access points and other active electronic infrastructure components within easy reach of local employees using commonly available ladders and stepladders. This facilitates easy removal and installation of these components should they require replacement. An example of this is shown in [Figure 11-7](#). In this case, a nationwide retailer has mandated that all electronic infrastructure components be accessible to store employees using the ten foot step-ladders commonly available at each store location. Here we see that the access point is mounted at 10 feet (for easy access) while the antennas are mounted at 15 feet. In cases such as this, the value specified for “AP Height” in [Figure 11-6](#) should reflect the height of the antennas and not the height of the access point.

Figure 11-7 Example of Different AP Heights



223850

**Note**

For a more detailed discussion of antennas, see:

[http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-antennas-accessories/product\\_data\\_sheet09186a008008883b.pdf](http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-antennas-accessories/product_data_sheet09186a008008883b.pdf).



## 802.11 Fundamentals

---

September 4, 2014

This part of the CVD discusses 802.11 Fundamentals, namely the role of beacons, probe requests, and probe responses.

Before discussing details about a location ready WLAN design, it is important to understand how an Access Point and wireless client start their initial communication. Once the wireless network is up and running, WLAN clients connect to an Access Point that provides it the best connection and data rate possible. Before a client connects to an AP, the client must first figure out to which AP it should connect.

802.11 WLANs consist of multiple elements and behaviors which make up the foundation of the 802.11 protocol. A key part of the protocol discovers the appropriate WLAN and establishes a connection with that WLAN. The primary components of this process are:

- Beacons—Used by the WLAN network to advertise its presence.
- Probes—Used by WLAN clients to find their networks.
- Authentication—An artifact from the original 802.11 standard.
- Association—Establishes the data link between an AP and a WLAN client.

## Beacons

Although beacons are regularly broadcast by an AP, the probe, authentication, and association frames are generally used only during the association and re-association process. In a CUWN network, the APs advertise their presence in the network by sending out Beacon frames, which includes the SSID and BSSID information. The Beacon frame also contains information about the supported rates, parameter sets that indicate channel number, security requirements (WEP or WPA, etc.), and optionally Traffic Indication Map (TIM) that APs can send periodically to poll stations that use power save mode but have data frames waiting for them at the Access Point. Typically APs transmit beacon frames every 100ms.



**Note**

Many WLAN security documents suggest that sending beacons without the service set identifier (SSID) is a security best practice that prevents potential hackers from learning the SSID of a WLAN network. All WLAN solutions offer this as an option. However given that the SSID can be easily discovered while sniffing a WLAN client during the association phase, this option has little security value. For a Cisco Connected Mobile Experience (CMX) Solution, SSID can be broadcast or hidden based on the use case. Generally Connect & Engage or Analytics services would presume that SSID be broadcast to get maximum results, it is not a must. For operational and client support issues, it is often better to allow the SSID to be broadcast. The SSID should be chosen to represent to the identity of the company/venue/mall

or the purpose of the WLAN, while at the same time being as unique as possible; the SSID should not give away the purpose or the owner of the WLAN. Creating long random strings as SSIDs is not recommended because this simply adds to the operations and maintenance overhead without an appreciable security improvement; a simple word is often the best choice. Common WLAN-related words should be avoided because there is no process or standard to prevent accidental or intentional SSID duplication.

The following is an 802.11 beacon example:

```
Type/Subtype: Beacon frame (8)
...
  Destination address: Broadcast (ff:ff:ff:ff:ff:ff)
  ...
  Sequence number: 2577IEEE 802.11 wireless LAN management frame
  ...
    SSID parameter set: "wpa1"
      Tag Number: 0 (SSID parameter set)
      Tag length: 4
      Tag interpretation: wpa1
    Supported Rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
      Tag Number: 1 (Supported Rates)
      Tag length: 8
      Tag interpretation: Supported rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
[Mbit/sec]
  ...
    Vendor Specific: WPA
      Tag Number: 221 (Vendor Specific)
      Tag length: 28
      Tag interpretation: WPA IE, type 1, version 1
      Tag interpretation: Multicast cipher suite: TKIP
      Tag interpretation: # of unicast cipher suites: 2
      Tag interpretation: Unicast cipher suite 1: TKIP
      Tag interpretation: # of auth key management suites: 1
      Tag interpretation: auth key management suite 1: WPA
      Tag interpretation: Not interpreted
  ...
```

## 802.11 Join Process—Association

Before an 802.11 client can send data over a WLAN network (Fast Roaming is an exception to this process, but is not discussed in this guide), it goes through the following three-stage process:

- 802.11 probing—802.11 networks make use of a number of options, but for an enterprise deployment the search for a specific network involves sending a probe request out on multiple channels that specifies the network name (SSID) and bit rates.
- 802.11 authentication—802.11 was originally developed with two authentication mechanisms. The first one, called “open authentication”, is fundamentally a NULL authentication where the client says “authenticate me” and the AP responds with “yes”. This is the mechanism used in almost all 802.11 deployments. Open authentication is the recommended choice for a Cisco Connected Mobile Experience (CMX) Analytics deployment, but not a must. Likewise it may not be necessary to have Open Authentication mechanism for Connect & Engage or Mobile App engagement. The second type of authentication, namely the WEP/WPA/WPA2, is a shared key mechanism that is widely used in home networks or small Wi-Fi deployments. If needed be, these authentication mechanisms may also be used in a CMX deployment where an open network is not desired, but analytics is still preferred (for example, a hotel or lobby environment). Enterprise authentication is achieved by using 802.1X/EAP authentication mechanisms and is not discussed as part of this solution guide.

- 802.11 association—This stage finalizes the security and bit rate options and establishes the data link between the WLAN client and the AP. If a client has joined a network and roams from one AP to another within the network, the association is called a re-association. The primary difference between an association and a re-association event is that a re-association frame sends the MAC address (BSSID) of the previous AP in its re-association request to provide roaming information to the extended WLAN network.

In conventional WLANs, APs advertise their presence by sending out beacon frames which include their SSID and BSSID. Prior to association, clients gather information about the APs by scanning the channels one by one either through passive scanning or active scanning. In passive scanning mode the client station moves the radio into each channel and waits to listen for beacons on the channel. The client station listens for beacons containing SSID that it may have already connected to before. If the client receives beacons from multiple APs for the same SSID, it attempts to connect to the AP with the best RSSI (receiver signal strength indicator).

The clients also perform active scanning, wherein the client stations send out probe request frames on each channel. These probe requests may contain SSID of a specific WLAN that the station is looking for or the probe requests can also look for “any” SSID to find out all the SSIDs in the proximity of the client. These are requests for APs to send out information about themselves. APs respond to Probe Requests with probe response frames, the contents of which are similar to Beacon frames. The APs operating on a particular channel responds back to probe request with a probe response with its SSID, supported rates, and security rates.

If a client station receives probe responses from multiple APs (and/or multiple SSIDs), the client station uses RSSI of the AP as a judge to connect to an AP with best signal strength.

The following shows a segment of a sample probe request, where the WLAN client sends out a request for a particular SSID (wpa1).

```
IEEE 802.11 wireless LAN management frame
  Tagged parameters (31 bytes)
    SSID parameter set: "wpa1"
    ...
    Supported Rates: 1.0(B) 2.0(B) 5.5 11.0 6.0 9.0 12.0 18.0
    ...
    Extended Supported Rates: 24.0 36.0 48.0 54.0
    ...
```

The following shows a portion of a sample probe response, where an AP using the specified SSID responds with supported rate and security properties for that WLAN SSID.

```
...
IEEE 802.11 wireless LAN management frame
...
  Tag Number: 1 (Supported Rates)
  Tag length: 8
  Tag interpretation: Supported rates: 1.0 2.0 5.5 11.0(B) 6.0 9.0 12.0 18.0
[Mbit/sec]
  ...
  Tag interpretation: WPA IE, type 1, version 1
  Tag interpretation: Multicast cipher suite: TKIP
  Tag interpretation: # of unicast cipher suites: 1
  Tag interpretation: Unicast cipher suite 1: TKIP
  Tag interpretation: # of auth key management suites: 1
  Tag interpretation: auth key management suite 1: WPA
  Tag interpretation: Not interpreted
...

```

**Note**

---

The user has to connect to an SSID irrespective of the type of authentication method used on the WLAN. Most of the client stations present the user with a list of SSIDs to join based on the best signal strength they receive. Once the user instructs a client to join a particular SSID, the client software usually stores this information in its cache; this info is then used by the client station to probe for the same SSID when it tries to connect again after a disassociation.

---

It is also fairly common for client station that is already associated with an AP to send probe requests every few seconds across other channels. The client station does this to maintain an updated list of known APs with best signal strength. When the client can no longer maintain a good connection with the AP, the client can roam to another AP with better signal strength. With newer improvements on the Cisco Wireless LAN Controllers, APs may force the client to disassociate so that the client can enter the 802.11 join phase again to connect to a better AP. Once the client station decides to join a particular SSID to the strongest AP, it follows through with the 802.11 authentication and 802.11 association phases to connect to the network.

**Note**

---

More detail on the 802.11 authentication frames and association frames are can be found at: [http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Mobility/secwlandg20/sw2dg/ch3\\_2\\_SPMb.html#wp1056095](http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Mobility/secwlandg20/sw2dg/ch3_2_SPMb.html#wp1056095).

---

This method of actively scanning by sending probe requests and receiving probe responses on different channels, irrespective of an on-going connection between an AP and client, makes it possible for a client and APs to know about each other constantly. The WLC and MSE make use of this information to locate a Wi-Fi client and form the basis for a location aware client design. [Chapter 13, “Location Fundamentals”](#) covers location awareness in more detail.



## Location Fundamentals

September 4, 2014

This part of the CVD discusses location fundamentals, including the definition of a location ready point, location currency, location accuracy, and location latency. We also discuss two methods of getting location from a client, i.e., the Probe RSSI method and the FastLocate method.


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

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

A point in a WLAN deployment is location ready if the following are all determined to be true:

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

Figure 13-1 Location Ready Point



The point in red in [Figure 13-1](#) may be thought of as a Wi-Fi client at a given location in a venue. In simplistic terms we can say that the MSE triangulates the location of a client by the RSSI at which client is heard at different APs. At a minimum for triangulation to work, three APs should hear the client at a RSSI level above -75 dBm. The MSE uses this information as well as device fingerprinting to calculate the location of the device.

The Cisco CMX solution gathers location of the client using two methods:

- [Probe RSSI](#)
- [FastLocate](#)

## Probe RSSI

The default way in which MSE computes the location of a client is via triangulation and finger printing of the device. The MSE triangulates the location of the client by the RSSI of the probe packets which are heard by different APs. From [802.11 Join Process—Association](#) in [Chapter 12, “802.11 Fundamentals,”](#) we know that Wi-Fi clients perform active scanning by sending out probe requests on each channel to look for a SSID to which to connect. The APs respond to the probe request with a probe response detailing characteristics of what they can serve to the client. The client in turn makes a decision to join an AP that provides the best signal strength. We also know that clients periodically perform active scanning in background, which helps them keep an updated list of Access Points with best signal strength to which to connect. When the client can no longer connect to AP, it uses the AP list stored to connect to another AP that gives it the best signal strength.

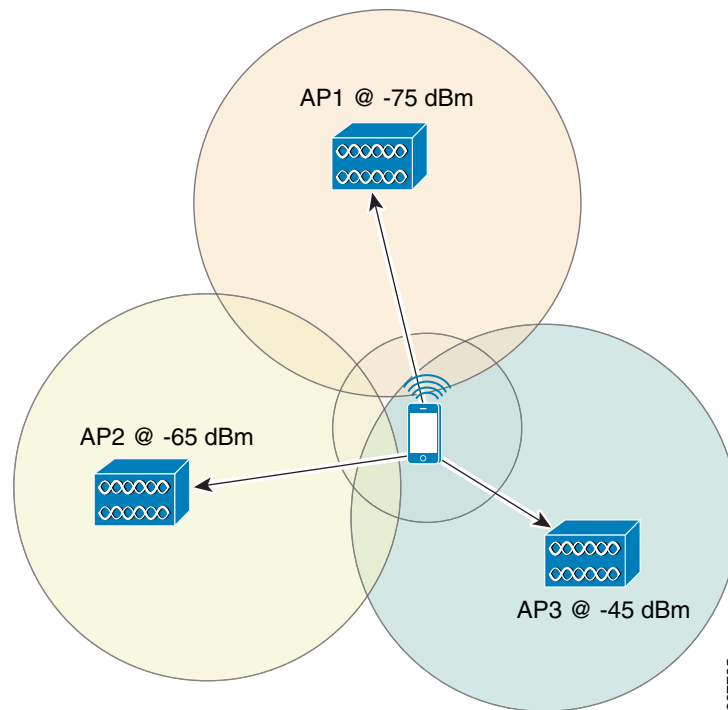
This back and forth exchange of Probe Requests and Probe Responses is also useful to MSE to calculate the location of the client. Just as the probe response is heard by the AP at a particular RSSI, the probe request of a client is heard at the AP at a particular RSSI. In [Figure 13-1](#), consider the red dot as a Wi-Fi client. The way the MSE would locate the client is:

- Client broadcasts probe requests on all channels.
- APs near the client respond back to the client with probe response. At the same time MSE creates a client entry with the RSSI at which different APs heard the client.
- MSE knows the location of all APs on the maps. Based on the location of APs on the maps and with triangulation, the MSE can approximately triangulate the location of the client based on which APs hear it.

For location of the client to be accurate, it is important that at a minimum at least three APs can hear the client. If the Wi-Fi client is heard by fewer APs, the accuracy of the location calculation can vary significantly because while probe RSSI tells the APs the strength at which they can hear the client, it does **not** tell the direction of the client where it is coming from.



**Figure 13-2** Location Readiness of Client



In [Figure 13-2](#), three APs can hear the client probe requests at different RSSI. However none of the APs can determine the direction in which the client is by itself. The MSE knows the location of APs on the Maps. Using the location of APs on the maps and RSSI at which client has been detected by APs, the MSE can triangulate the approximate location of the client.


However it is wrong to assume that because three or more APs can hear a client that the client “location” can be accurately tracked. A good location ready deployment is a deployment that not only determines the client existence, but also the location of the client with reasonable accuracy (the MSE provides a 90% confidence level that a client will be within 10 meters of its actual location). To get to that level of accuracy, it is very important to understand Access Point placement and how to deploy a WLAN that is location aware and ready. [Chapter 14, “Pre-Deployment Radio Frequency Site Survey”](#) goes into more detail about AP placement and RF planning.

## Location Latency

While the Probe RSSI method was a good way to calculate the location of the client, it has some limitations with current mobile devices, tablets, and smartphones. Traditionally the Wi-Fi clients transmitted probe requests every 10ms seconds which provided MSE with enough probe requests and RSSI information to calculate the location of the client. Transitions in technology have meant that the latest clients do not nearly probe for Wi-Fi networks as frequently as they used to. iOS-based devices such as iPhones, iPads, etc. probe much less frequently (up to 30 seconds between probes instead of 1-2 seconds previously). Moreover Wi-Fi clients in general do not tend to probe for Wi-Fi networks if they are in motion. Some clients may probe more when not associated versus when they are actively associated with a SSID. Some devices may send probes regularly at intervals, while some devices may do probes in burst. To save on battery and power, devices may not probe frequently. All these issues can have a negative effect on location accuracy. For more details on refresh rates of devices, see [Chapter 6, “CMX Additional Considerations.”](#)

Figure 13-3 illustrates how long it can take for a client's location to be determined with the Probe RSSI method.

**Figure 13-3 Location Latency**



- Wi-Fi Client probe (5 seconds to 5 mins depending on client).
- AP radio reports strength of client in dB every 500ms.
- AP groups messages and sends them to WLC via CAPWAP every 2 seconds.
- WLC aggregates and sends RSSI report to MSE via NMSP (min 2 seconds, default at 4 seconds).
- MSE Aggregates data for 5 seconds, calculates location (sends push and batches db write).

As seen above, in the best case scenario using the probe RSSI method, determining the initial location of a client that sends probe requests every 30 can be calculated as:

Initial Client probe + 2.5 seconds (AP reports & sends) + 2 seconds (best case WLC ->MSE transfer) + 5 seconds (MSE aggregation window) = 8.5 seconds.

An update to the client's location however if the client is sending probes every 30 seconds would be:

Client probe (30seconds) + 8.5 seconds = 38.5 seconds.

Note that not all clients send probe requests at 30 seconds; they may send at lower or higher rates. Also once the client sends the probes, the client waits in the channel for 10ms for the probe response. Again this value of 10ms is not mandated and a client Wi-Fi station may choose to wait longer. However as we are only concerned with RSSI of the probe request, we take that into account to see how quickly we can determine the location.

The problem is clear from above. As the probe requests get slower and devices tend to move, they send out infrequent probe requests. These infrequent probe requests mean no updates to the MSE, so a client location may not be updated on the map because the MSE does not have enough information to update the location of the client. This may result in unsatisfactory experience for a CMX deployment where location is the key to provide services to client and also for location analytics.

To improve location updates, starting with CUWN 8.0 release, MSE can now use FastLocate or Data packet RSSI to get faster updates on RSSI. The FastLocate or Data RSSI method is detailed below.

# FastLocate


**Note**

The FastLocate feature is only available with addition of the WSM module.

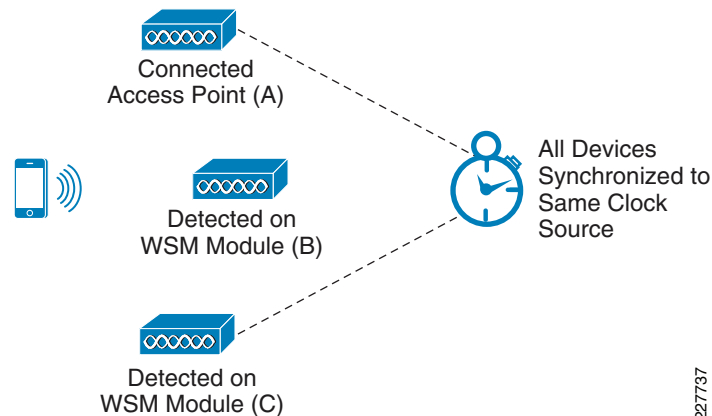
FastLocate with addition of the WSM module provides a way to mitigate the low frequency of probe requests. Note that FastLocate does not improve location accuracy; rather, FastLocate provides a way of getting faster updates about a client and signal it can be heard. This increase in frequency means that location updates for a client are faster, even if the client is in motion. FastLocate supplements Mobile Probing with All Network Packets for Higher Location Resolution. Some of its advantages include:

- Initiated by network—Available more frequently
- Provides more data points to accurately represent end user activity
- Device agnostic—Consistent across devices and works even when device is sleeping
- Has minimal impact on battery life of mobile device

The idea behind FastLocate is to use the WSM module to gather RSSI information from data packets coming in from client devices and pass it on to the MSE to calculate location. The advantage of using data packet RSSI is that it is more likely that a Wi-Fi client is sending data after it is associated with a network than not. An application that uses CMX infrastructure may also be configured to send data packets proactively to enable better location updates, however it is not necessary.

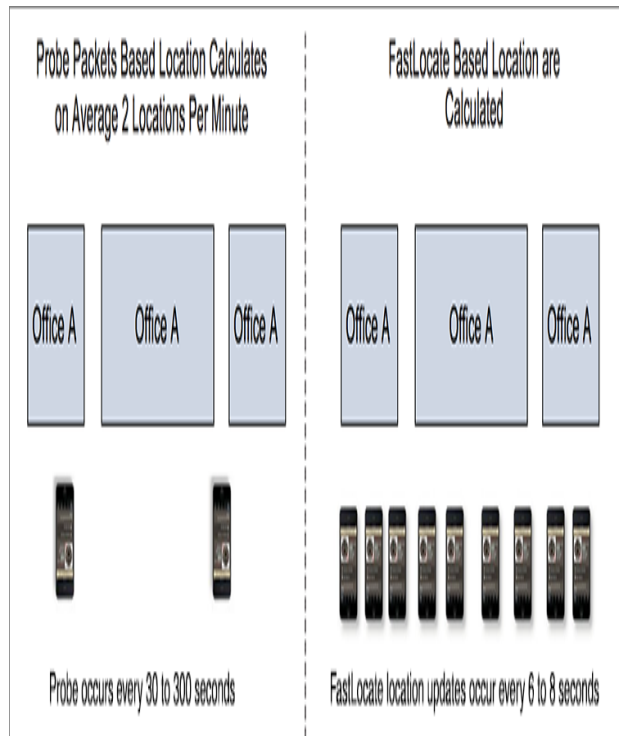
Figure 13-4 illustrates how FastLocate works.

**Figure 13-4** *Fast Locate Operation*



- All devices (Access Points, WLC, MSE) are synchronized to the same clock source. This is very important requirement for FastLocate.
- Associated device send packets for regular data access only to connected access points (A).
- Other AP (B,C) that “hear” that MAC address talking to associated AP can report on signal strength to WLC/MSE.
- Additional smart techniques are used for quiet devices (for example, an App on the device can send data packets to AP periodically).
- The WLC aggregates the data as before and sends this information back to MSE.
- MSE locates the client as before, but now has more RSSI measurements more frequently, and so MSE is able to better update the location.

**Figure 13-5 Probe RSSI versus Fast Locate**



Because the WSM module can scan channels independently and on demand, FastLocate enables faster location calculations. FastLocate- location updates occur almost every 6-8 seconds in contrast to a probe RSSI method where the device may send a probe anywhere from 30 to 300 seconds. Clearly there is a huge advantage to using FastLocate in a CMX deployment.

However frequent location updates should not be confused with location accuracy as the location accuracy of a client is different from location updates. Irrespective of whether probe RSSI or FastLocate are used, an understanding of Location Accuracy and Currency is important while designing a location aware CMX deployment.

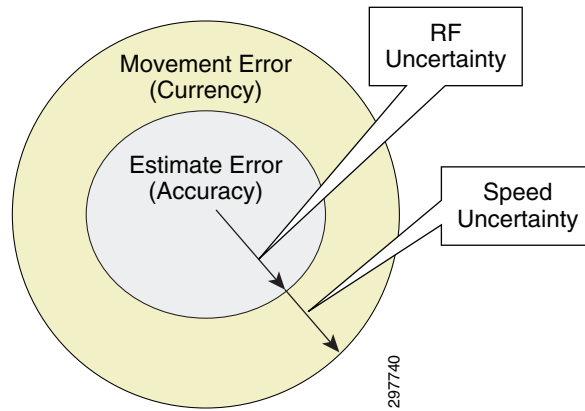
## Location Accuracy and Currency

Location accuracy of a given client is the magnitude of error between the actual position and predicted position. MSE relies on RSSI collected from various access points to calculate location. Factors like AP height, AP density require careful planning. RF reflection and multipath can cause large variations in RSSI even at small distances. Mobile clients' transmit power can change over a short period of time even when client is stationary at the test spot. Client software drivers that control the transmit power can cause differences between mobile devices and in some cases within different software versions of the same device type. Also different mobile devices can vary in their antenna placements. For these reasons, sometimes location calculated for two mobile devices at the same location can vary. In practice, to analyze the location accuracy of a system, statistical tools are used.

When evaluating location accuracy, multiple test points throughout the venue are chosen. Several tests are run with various types of mobile clients at designated spots to obtain a statistically significant sample set of data. Clients are stationary at the test spot. Location error from these tests are analyzed using statistical tool such as a cumulative distribution function (CDF). 50 and 90 percentile location errors are derived from the CDF and are important metrics to understand before planning location based services.

Accuracy (error in prediction) is distinct from currency (last time of prediction) which is a function of the client device probing frequency and not the Wi-Fi infrastructure directly (following the guidelines herein will maximize currency).

**Figure 13-6 Accuracy and Currency**



From [Figure 13-6](#) we see that the total location error is the sum of the initial position estimation error (accuracy) [as a result of the uncertainty in exact RF propagation loss between the client and the AP] plus the movement error (currency) as a result of the uncertainty in speed of the user (since the client device reports position periodically).





# Pre-Deployment Radio Frequency Site Survey

---

September 4, 2014

This part of the CVD discusses the pre-deployment RF site survey. A good Cisco WLAN deployment is dependent on a good RF design, which includes doing a thorough site survey of the location, determining the best location for access points, making the right channel plans, planning for AP capacity, and lastly performing a regular post deployment RF site survey. The chapter discusses pre-deployment site survey topics.

## Pre-deployment RF Site Survey

A pre-deployment site survey is one of the most important tasks that should be conducted before deploying a Cisco Wireless LAN. Pre-deployment RF site surveys involve physically inspecting the location where the WLAN needs to be deployed, studying building properties to plan for multipath, signal attenuation, etc., and verifying that Access Points can be installed in a location that provides maximum coverage with good signal strength to users. A typical pre-deployment RF site survey should involve discussion and planning around the topics discussed in the following sections.

## Physical Site Survey

The purpose of a physical site survey is to establish a firm view on planning for a WLAN RF site. Typical items to consider in a physical site survey include gathering knowledge not just about existing WLANs at the site, but also considering building types, coverage requirements, and restrictions. The sections below describe the typical data that should be collected.

## Location Assessment

- Assess building type and materials used—As previously mentioned, building type and material play a big part in attenuating signal, which in turn means planning for coverage.
- Anticipate difficult zones—Coverage in stair cases or atrium like environments.
- Areas where full coverage and full performance are needed—Key locations where full wireless coverage and performance are required should be noted in advance.
- Areas where location is important and needed—Certain areas can be marked as a no Wi-Fi zone or exclusion zone in which a customer may not be interested. These should be marked on map.

- Areas where RF free zones exist—Unexpected roaming paths that a user may take in the building? This information is very useful in later parts when analyzing analytics data as well as planning for better mobility.
- Obtain location maps for RF surveys later—It is important to obtain updated maps for locations where CMX services should be deployed. The maps should be accurate as these maps are for AP placement and capacity planning. The maps also help in determining predictive RF site survey.
- Outdoor readiness—While the Cisco CMX solution guide does not cover outdoor deployments, in the interest of completeness it is important to consider the following for an outdoor environment. Most of the guidelines that apply to indoor deployment also effectively apply to outdoor deployments. Other guidelines to consider may include outdoor casing for APs, power requirements, aesthetic constraint on outdoor requirement, and mesh network requirements if any. Outdoor networks are more susceptible to RF bleed from other networks than indoor networks. A survey of existing 802.11 coverage should be done to analyze how much RF bleed occurs and what can be done. Safety of installation of Access Points with other RF equipment in area should also be considered.

**Note**


---

Presence Analytics can be used for outdoor environments where a single AP in a store or few APs can be used for a presence site. However this guide does not cover this scenario in this release.

---

- Location where Access Points can be deployed—Network drops and Access Point instability should also be considered. A general rule of thumb is that Access Points should be installed less than 70ft from each other to ensure good location coverage. Is the location capable of providing AP drops where needed? Potential sources of interferences should be noted, such as static analog video cameras and microwave ovens.
- High Density—Is this location going to be a high density location. High density is defined as a location that has a large number of clients or a network that aims to provide high level of bandwidth to all its clients. [Chapter 18, “Capacity Planning and High Density”](#) goes into more detail about high density deployment recommendations.

## Business Needs of WLAN

- Critical Services—What services does this business want to provide to its users? For a typical Cisco Connected Mobile Experience it is important that CMX can detect, connect, and engage its users effectively.
- Long term Network View—Planning for future needs is always better than planning for the present. With the explosion of smartphones and mobile devices in markets, more and more devices are constantly on the Wi-Fi network. The Internet of Everything (IoE) is enabling machine-to-machine and machine-to-network communication via Wi-Fi too. It is not enough to plan for smartphones and laptops alone, but also for other potential IoE devices that may use Wi-Fi.
- Guarding against potential RF explosion—As machine-to-machine and various devices proliferate on the network, RF spectrum will be at a premium. Innovations like CleanAir enable networks to heal themselves and operate better.
- Scale of the planning—One building versus multiple buildings, malls, hospital size, etc. dictates the scale of the installation. Is this a mission critical network or a guest network (like malls, etc).



- Expected audience of the network—Are your users going to be mostly smartphone/tablet users (such as in retail establishments and malls) or users with voice applications running on laptops. While there is no way to predict exactly how many user devices will be on a network, a good ball park estimate should be used while designing the network for capacity.
- 802.11n/ac readiness and expectations—It is important to study the physical aspects of location and determine 802.11n/ac readiness of the deployment. While 802.11n/ac provides higher speeds, they also come with the price of lost range. A good multipath environment is good for an 802.11n/ac deployment, while a hallway deployment may need more Access Points because of lesser multipath. It is recommended to keep the deployment consistent by using Access Points of the same type throughout the location.



---

**Note** Note: 802.11ac does not have any impact on location accuracy or calculation of clients.

---

## Constraints on Deployment

- RF Free Zones—Is there a need for RF free zones in building where no coverage is desired?
- DFS and radar avoidance requirement—Know in advance the DFS and Radar requirements for deployment. Different countries have different regulatory domains, so have a list ready of available channels and power levels that can be deployed ready in advance.
- Aesthetic design requirements—Aesthetics should be considered and planned for. For example, a hospital may not want any external antenna to be visible or a mall may want to hide all Access Points.



**Note**

---

[Chapter 11, “Antenna Fundamentals”](#) discusses various antenna options when deploying a WLAN solution. Antennas that are a best match for the venue and aesthetics should be chosen.

---

## Budgeting

- Predictive surveys for simple budgeting—Based on the above information, a simple budgeting estimate should be done to avoid any potential surprises in deployment costs later.
- Plan for cabling, power drops, and power requirements—Plan for obtaining cabling and power drops. You may discover later that an ideal location for AP does not have power drops. Can they be added later?
- Thorough versus sample area survey—Sometimes it may not be necessary to evaluate the entire location. A sample area might be enough to extrapolate information. This is usually sufficient in a location where it is determined that there are very few external factors that might hinder deployment. Another example is a huge mall deployment where it might be a more cost effective use of time to survey the edges of the mall for existing RF rather than center of mall.

## Existing 802.11 Surveys

- It is a good idea to capture existing 802.11 networks and their properties at the location where CMX needs to be deployed. Because Wi-Fi and other technologies are hugely prevalent today, it is a good investment of time to have a pre-study of existing RF done. Use tools like Metageek site survey to record existing RF at the location. In a big location it is advisable to move through the entire location and take a RF base reading.
- Plan for persistent non-movable interferers—Plan for existing RF deployments that can potentially interfere with your CMX deployment. (analog security cameras, microwave ovens, etc.). Existing outdoor bridges like Motorola Canopy wireless systems may operate on the same spectrum as other 5Ghz systems. Since these antennas usually cannot be moved, care should be taken to take note of where these systems are and plan APs around them.
- Capture current state of Wi-Fi network—How many SSIDs exist already? What is their signal strength? Are there potential Wi-Fi networks around the location that are strong enough to interfere with a CMX deployment? What can be done to mitigate such potential sources of interference?
- Antenna Evaluation—Evaluate using external antennas instead of internal antennas to ensure good coverage based on above parameters? Will use of external antennas be an aesthetic constraint? Can directional antennas be used in corner of the building to direct more coverage into the building rather than outside of it? In a large environment this may be useful.

## Use Case Example

As an example to illustrate these considerations, consider a deployment using a floor space similar to that discussed for location analytics in [Chapter 7, “CMX Use Case Stories.”](#) Using the maps and doing a physical site survey of the location, we can determine in advance many details that can be used before doing an actual RF site survey and planning.

[Figure 14-1](#) illustrates an example of how a physical site survey results in information gathered for a RF site survey and RF planning.

Figure 14-1 Use Case Example Layout







# Access Point Placement and Separation

---

September 4, 2014

This chapter discusses AP placement and AP capacity planning, including core concepts regarding the distance between APs in a network and its impact on location, data, and voice.


## Access Point Separation

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

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

The techniques incorporated in the location-aware Cisco UWN to localize tracked devices operate most effectively when RSSI and distance are seen to possess a clearly monotonic relationship. To better understand what this means, we examine a simulated plot of a tracked device's detected RSSI as the distance between it and a detecting access point is increased. While the relationship between RSSI and distance varies depending on different combinations of antenna, antenna height, and environmental characteristics, the graph shown in [Figure 15-1](#) is for an access point mounted at approximately twelve feet elevation can be used to better understand the concepts discussed here.

**Figure 15-1** Relationship between RSSI and Distance



In [Figure 15-1](#), we see that beginning at some point a fairly near the access point and ranging to another point c further in distance, the two variables exhibit a strict monotonically decreasing relationship (as distance between the tracked device and the access point increases, the RSSI at which the access point detects the device is shown to decrease). Between point a and another point b, the amount of change in RSSI (dBm) that occurs per-unit change in distance (feet) is highly consistent, approximately -5 dBm per 20-foot change in distance. This results in the slope of the graph between points a and b being fairly steep. As the distance continues to increase beyond point b, the slope of our graph begins to diminish and the level of RSSI differentiation decreases, providing increasingly less differentiation in received signal strength per-unit change in distance. Note that the slope of the graph between points b and c is not nearly as steep as it is between points a and b. As distance begins to significantly exceed point b in this example, the slope of the graph diminishes even further. This greatly reduced slope and steepness results in a decreased level of differentiation in signal level with increasing distance. When this occurs at extended distances, it becomes more difficult to accurately predict changes in distance based on detected changes in RSSI (lateration).

The risk of this lack of RSSI differentiation having a significant impact on location accuracy can be reduced if steps are taken to avoid areas of the RSSI versus distance curve where this phenomenon is known to exist most prominently. In general, for access points deployed indoors at antenna heights of 20 feet or less, this can be achieved if the range of any point on the floor to at least three detecting access points on that floor (one in each of at least three of the four quadrants surrounding it) is maintained within approximately 70 feet in an indoor environment. This is a general recommendation that is intended to assist designers in avoiding situations where excessive inter-access point distance may be a contributing factor to location inaccuracy. As shown in [Figure 15-1](#), diminished RSSI differentiation with increasing distance is a gradually increasing phenomenon, therefore a degree of flexibility is implied in this recommendation.

In practice, in addition to being conducive to good location accuracy, this recommendation applies well to deployments where location tracking is deployed in conjunction with other WLAN applications (such as voice and high speed data) in accordance with current recommended best practices. This is especially true for environments where the expected path loss exponent is 3.5 (walled office environment) or higher, as the required inter-access point spacing tends to generally fall within this range. In addition to

the potential effects of a lack of RSSI differentiation at distance extremes, inter-access point distances significantly greater than 70 to 80 feet can make it more challenging to satisfy the best practice signal strength and overlap requirements in environments with high path loss.

At ranges closer than point a in our example, propagation anomalies that are due to the elevation pattern of the chosen antenna, the antenna's installation height, and the current physical location of the tracked device can potentially combine to degrade monotonicity. As a result, RSSI cannot be depended on as a reliable predictor of distance in this part of the curve, since it may be possible that more than one equally likely value for distance exists at a particular detected RSSI level. [Figure 15-2](#) illustrates an example of close-range non-monotonicity, depicting how a tracked device's RSSI reading of -40dBm can be associated with three different distances (5, 7, and 12 feet) from the access point antenna when operating in this close-range non-monotonic region of the RSSI versus distance graph. This behavior is typically the result of a variation in an overhead antenna's propagation pattern as a device approaches it begins to venture into the area almost directly beneath it. Obviously, these effects vary depending on the propagation pattern of the specific antennas used and their installation height above the area where tracked devices is located. However, the lesson to be learned from this is that although increased access point density can often be conducive to better location accuracy, the effect is not without its limits.

**Figure 15-2** Close Range Non-Monotonicity



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


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

location fidelity. One way to accomplish this for co-located access points in a lecture hall, for example, would be to place the access points on different walls and perhaps the ceiling as well with appropriate inter-access point spacing.

In general then, most indoor location tracking deployments with access point antennas installed at heights of between ten and twenty feet can be well served with an inter-access point spacing of between 40 and 70 feet, especially when combined with the signal threshold and access point placement recommendations suggested in the preceding sections of this document. In some cases however, inter-access point spacing below 40 feet may be necessary to satisfy the requirements of some applications for high signal strength thresholds, especially in environments where high path loss is present. An example of this might be a voice application deployed in such an environment (for example, a path loss exponent of 4.0 where a high degree of environmental clutter is present). Best practices for Location Aware deployments for CMX suggest a minimum signal level of -67dBm, 20% inter-cell overlap, and signal to noise ratio of 25 dB for 802.11n in this type of situation. Applying these requirements mathematically, we calculate an estimated cell size of 24 feet and an inter-access point spacing of 33 feet. In this case, to deploy our voice application in accordance with recommended best practices, the inter-access point spacing should be reduced below the general guideline of 40 feet. Note that good location accuracy is achievable at inter-access point ranges below 40 feet, provided that the access point spacing is not decreased so much that the negative effects of close range non-monotonicity come into play. Generally, this should not be an issue if the inter-access point distances are above 25 to 28 feet when using low gain, omni-directional antennas mounted at an installation height of approximately 10 feet in an indoor environment.

Figure 15-4 shows an example of location aware AP placement that illustrates access point placement and inter-access point spacing, offering a foundation for a location-aware design. The environment in Figure 15-4 consists of drywall offices and cubicle office spaces with a total space of approximately 275 feet by 159 feet. Taking into consideration the location tracking requirement for illustrative purposes only, our inter-access point linear-spacing recommendations of 40 to 70 feet suggests approximately 22 location-aware access points as an initial estimate. Incorporating the placement strategies made in preceding sections, interior, perimeter, and corner access points are placed to facilitate multi-lateration and establish a clearly delineated convex hull around the floor.


**Figure 15-3 Inter-AP Distance**



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



Figure 15-4 Example of Location Aware AP Placement



The Cisco Prime Infrastructure includes a planning tool that allows designers to model “what-if” design scenarios. This is a predictive modeling tool that is used on a per-floor basis to provide initial guidance on access point placement, as well as an interactive representation of predicted access point signal strength and data rate information. It can be safely used without impacting any actual deployment of access points that may already be in service. It is however possible that Cisco Prime Infrastructure may not be available at a location when the deployment is being designed. Other RF planning tools like Ekahau can also be used to plan for RF in advance with scenarios in mind.

**Note**

The Cisco Prime Infrastructure RF Planning tool and Ekahau RF Survey planning tools are both discussed in more detail in [Chapter 16, “Predictive Radio Frequency Planning.”](#)

## AP Placement

Proper placement of access points is one of several best practices that should be adhered to in order to unleash the full performance potential of the location-aware Cisco Unified Wireless Network. In many existing enterprise LANs, access points are distributed mainly throughout interior spaces, providing service to the surrounding work areas. These access point locations have been selected traditionally on the basis of coverage, WLAN bandwidth, channel reuse, cell-to-cell overlap, security, aesthetics, and deployment feasibility. In a location-aware WLAN design, the requirements of underlying data and voice applications should be combined with the requirements for good location fidelity. Depending on the particular site, the requirements of the location-aware Cisco UWN are flexible enough such that the

addition of location tracking to voice installations already designed in accordance with Cisco best practices, for example, may not require extensive reworking. Rather, infrastructure already deployed in accordance with accepted voice best practices can often be augmented such that location tracking best practice requirements are met as well (such as perimeter and corner access point placement, for example) depending on the characteristics of the areas involved.


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

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

Figure 15-5 illustrates the concept of a convex hull. Assume the set of access point locations is denoted by the black dots, which we refer to as set  $S$ . The convex hull of set  $S$ , or  $\text{Hull}(S)$ , is figuratively represented as an elastic band (shown by the blue line) that is stretched and allowed to snap over the outermost members of the set (which in this case represents perimeter and corner access points).

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

**Figure 15-5** Convex Hull Set of Points



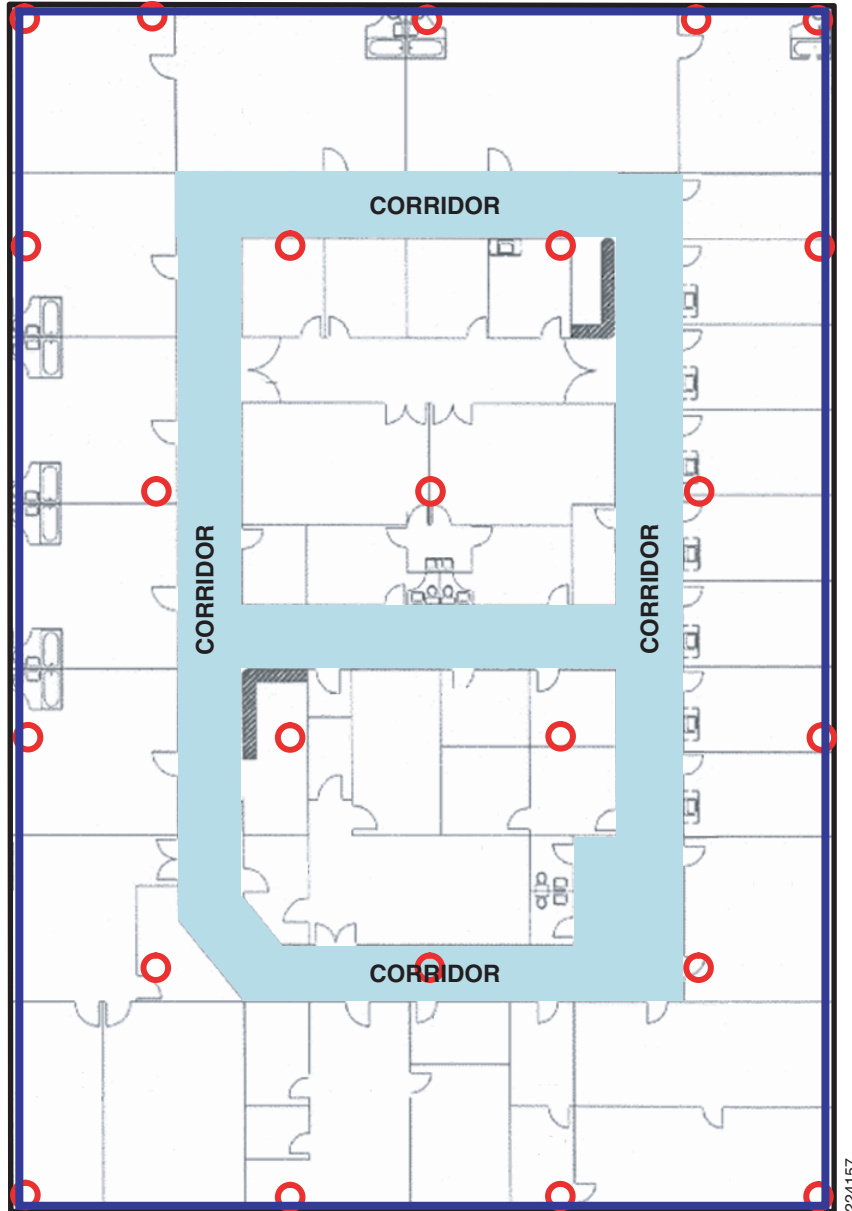
To ensure proper convex hull establishment around the set of location data points possessing high potential for good accuracy, access points should be placed in each corner of the floor as well as along the floor perimeter between corners. Inter-access point separation along the perimeter should be in

accordance with the general access point separation guidelines (described in [Chapter 16, “Predictive Radio Frequency Planning”](#)). The designer may reduce this spacing if necessary in order for these access points to participate in the provisioning of voice or data service to the floor.

## Proper Access Point Placement

[Figure 15-6](#) shows an example in which these concepts are applied to a type of floor plan found in many enterprises (that of rooms or offices contained by and surrounding an interior corridor). In this case, the area in which we desire to locate clients is the entire floor. In [Figure 15-6](#), note that the access points located towards the center of the floor are complemented by those that have been placed along the perimeter, but slightly inside. As is the case in most proper location-aware designs, the set of location data points possessing the highest potential for good location accuracy is contained within the convex hull, which in [Figure 15-6](#) is represented by the blue rectangle and encompasses the entire floor.

**Figure 15-6** Proper Access Point Perimeter Placement




In some cases, customer preferences or deployment restrictions may factor into the access point placement decision and the placement of access points at the floor perimeter may be restricted in one way or another.

Also it is not entirely necessary to place the APs right at the perimeter. In fact the APs can be slight staggered inside the building to provide better RF coverage as well as minimize RF wastage outside the building or necessary area.

# Improper Access Point Placement

Figure 15-7 illustrates an example of a less-than-desirable situation where the placement of access points has been restricted to hallway corridors and administrative/storage facilities located within the areas encircled by the corridors. For aesthetic reasons, facilities management has decided that access points will not be placed within any of the executive offices or conference rooms located between the hallway corridors and the physical perimeter. Because of these restrictions, our convex hull now lies at the outside edge of the corridor (indicated by the blue rectangle) and not at the true physical perimeter of the floor.

Figure 15-7 Improper Access Point Perimeter Placement



Given what we know about the distribution of location errors when operating outside the convex hull, it is logical to expect that location accuracy will not be as good in the offices and rooms located there. These areas of potentially lower accuracy are highlighted in red in [Figure 15-7](#).

With our recommendation of establishing the convex hull at the true floor physical perimeter notwithstanding, in practice the difference in location error rate between points located within the convex hull and outside it may be tolerable in some situations. These might include situations where such areas extend beyond the office perimeter for only a short distance (for example, small 10x10 foot rooms lining the walls of a corridor). For example, looking at the areas highlighted in red in [Figure 15-7](#), the potential increase in location error would be less in the smaller offices located at the right side of the floor plan than in any other affected area. Depending on magnitude, the effect of operation outside the convex hull will likely be the least. In contrast, the areas at the bottom of the floor plan, with larger offices and multiple wall partitions, would be potentially effected to a significantly higher degree.

In cases where access point placement in perimeter offices and conference rooms is restricted due to aesthetic concerns, a potential compromise may be possible using a very low profile antenna along with access point mounting in a plenum-rated enclosure (where permitted by local codes). This would offer the ability to mount access points at the proper perimeter and corner locations (thereby avoiding the quandary described in [Figure 15-7](#), but with a minimal visible footprint to the casual observer.

## Getting Around Placement

As mentioned earlier, the floor plans shown in [Figure 15-6](#) and [Figure 15-7](#) are commonplace, but by no means exclusive. For example, some modern building designs may possess hallway corridors that are located directly alongside the actual floor and building perimeter, typically allowing a panoramic view of campus environs as visitors move about between offices and conference rooms. In this case, all offices and conference facilities are located within the area between the corridors and the center of the floor. [Figure 15-8](#) provides an illustration of such a floor plan. Note that with this floor layout, placement along the outer edge of the hallway corridor places the access points along the actual physical perimeter, by default.

Figure 15-8 Perimeter Corridor Placement Plan

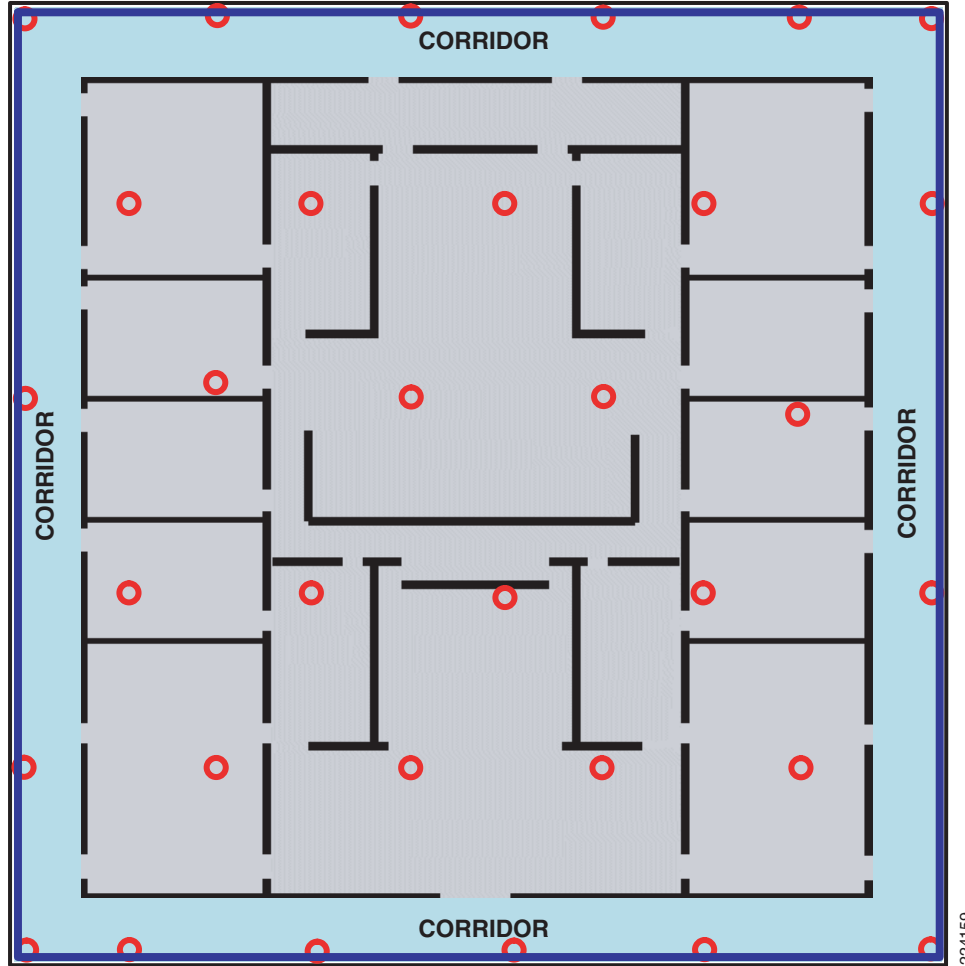
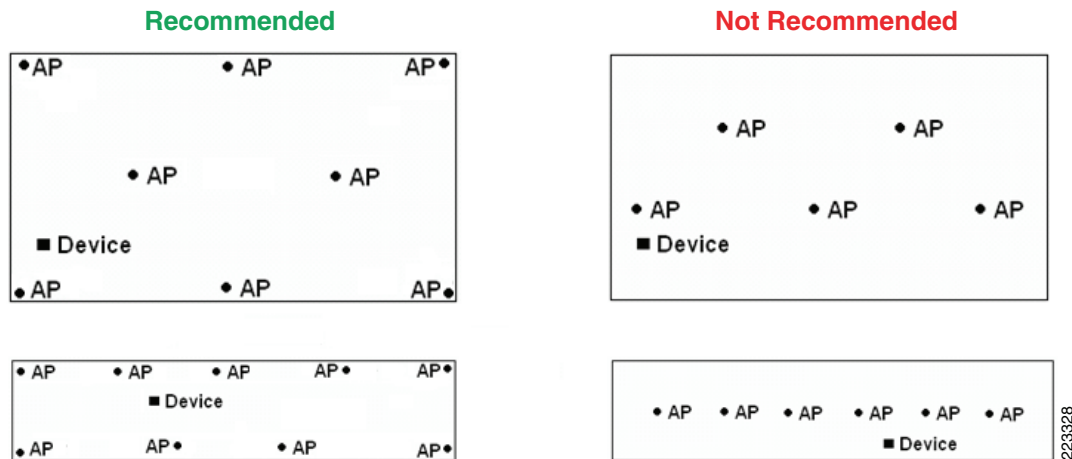


Figure 15-8 provides simple illustrations summarizing the access point placement concepts discussed in this section so far. Note that designs that make use of only clustered or straight-line access point placement should be augmented or redesigned in favor of those that combine center access point placement with perimeter and corner placement.


# Recommend Access Point Placement

Figure 15-9 Recommended Access Point Placement



If possible, mount antennas such that they have an unencumbered 360° view of all areas around them without being blocked at close range by large objects. For example, if possible, avoid placing access point antennas directly against large objects such as steel columns, as illustrated in [Figure 15-10](#). One option is to mount the access point along with its antennas to a ceiling location (provided that this allows an acceptable mounting height). Another option is to use short, low loss cable extension to allow separation between antennas and such obstructions.

Figure 15-10 AP Mounted Directly to Steel Column







# Predictive Radio Frequency Planning

September 4, 2014

This chapter discusses predictive RF planning that should be undertaken after a pre-deployment RF site survey is completed and two tools to perform RF planning, the Cisco Prime Infrastructure RF Planner tool and the Ekahau Site Survey tool.



**Note**

For smaller presence sites that involve only one or a few APs, a thorough predictive RF plan may not be required. However it is still necessary to do a limited site survey to understand the RF characteristic of the location before deploying APs.

After the physical site survey is done and analyzed, it is time to conduct an actual RF site survey to assess the location where CMX is to be deployed and determine an effective RF plan that requires minimum or no physical changes after it has been deployed. This is achieved through:

- Predictive planning through RF tools like Ekahau or Prime Infrastructure.
- Coverage check through a single AP in the location for RF characteristics.
- Post-deployment checks to ensure that RF coverage is indeed as determined.

Once a pre-site survey has been done and all the necessary information has been obtained, tools like the RF planning tool in Prime Infrastructure or Ekahau Wi-Fi planner should be used to do predictive RF planning. Perform as much background work as possible before doing predictive RF planning. Predictive RF planning does not guarantee that the network will behave exactly as the tool predicts, but it will generally have an error margin of 15-20% based on the background work done. For example, if a WLAN RF plan is done without any consideration given to an existing WLAN that already exists, then the RF plan, though theoretically accurate, might still not perform well. It is for this reason that after a deployment has been done, post RF site surveys are conducted to determine if coverage holes exist and whether they can be mitigated by either increasing power levels, adding access points, etc.



**Note**

Cisco does not recommend a particular RF planning tool. Tool from other vendors might achieve the same results, but for the purposes of this CMX CVD we use the Ekahau Wi-Fi planner and RF planning tool discussed in detail below.

Before starting predictive RF site planning, have the following information available:

- Maps of the location where CMX is to be deployed.
- Points on maps where access points can or cannot be installed.
- Points and areas on maps where full coverage is expected.

- Type of walls and materials used in the walls.
- If using external antennas, the type of external antenna used.

## Cisco Prime Infrastructure RF Planning Tool

The RF Planning tool can be used to add wall attenuation information to floor maps. Wall information added via the Map Editor does not affect access point placement or location designs, however it will be used by the planning tool when displaying predicted RF coverage maps for planned access points.

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

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


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

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

More complex designs containing totally enclosed interior voids (for example, a building with a fully enclosed interior atrium as shown in [Figure 16-1](#) with the perimeter of the building shown by a red outline) may not lend themselves well to automatic access point placement. The planning tool does not currently allow the exclusion of zones into which access point placement should not occur. Note in [Figure 16-1](#) the placement of access points 2, 4, 9, and 24 in the atrium area (indicated by the blue

outline). The placement of these access points in this area is incorrect, since the floor map is for the building's third floor. This should be corrected by manual intervention and moving the access points into correct locations or eliminating them entirely if not necessary.

**Figure 16-1** Example of a Floor Plan with Fully Enclosed Interior Atrium

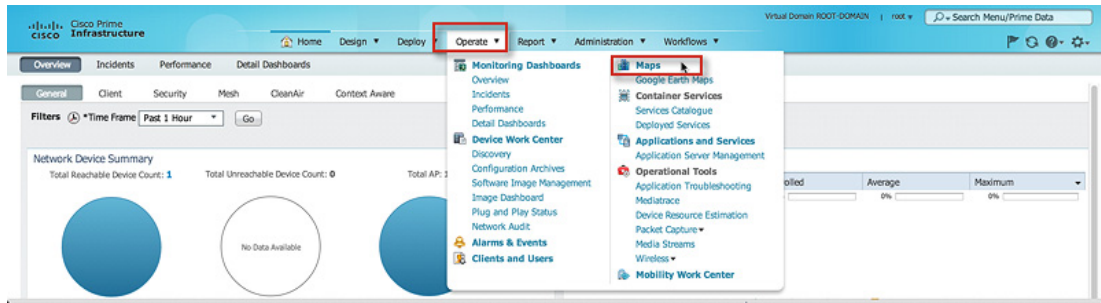


223333

To perform predictive RF Planning through RF Planner tool through Prime Infrastructure:

- 
- Step 1** Create and import maps into Prime Infrastructure.
  - Step 2** Click **Operate > Maps** to go to the Maps View.

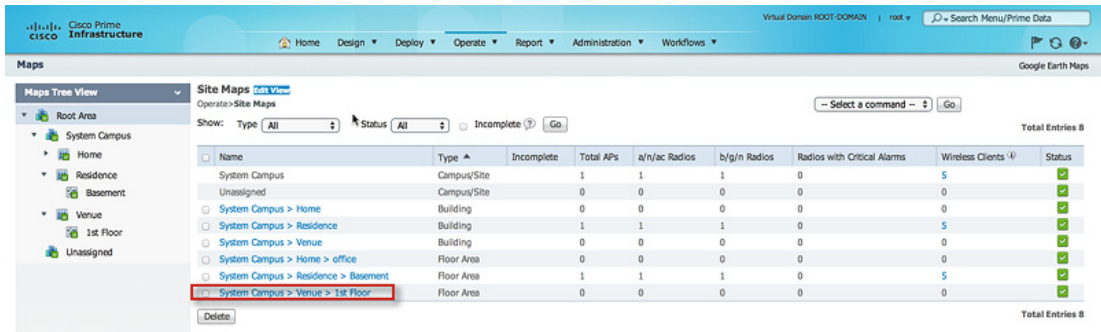
Figure 16-2 Maps View from PI



297743

- Step 3** Select the Floor for which you want to create an RF plan (this assumes that you have already uploaded a map to the floor).

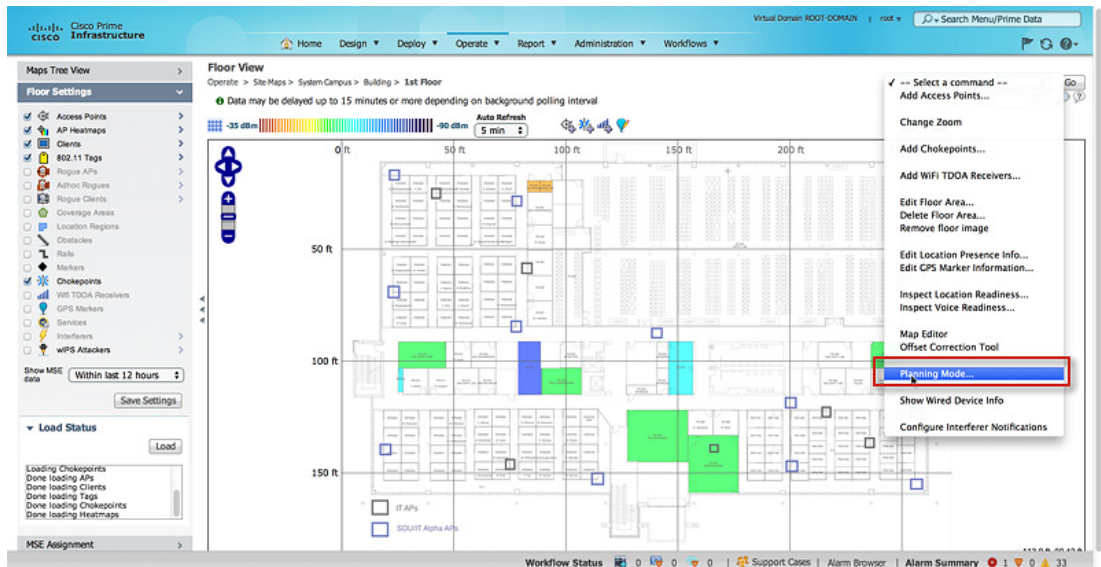
Figure 16-3 Select Floor for RF Plan



297744

- Step 4** The Floor view should display. In the top right “Select a command” drop-down menu, select “Planning Mode”. A new window is displayed.

Figure 16-4 Select Planning Mode

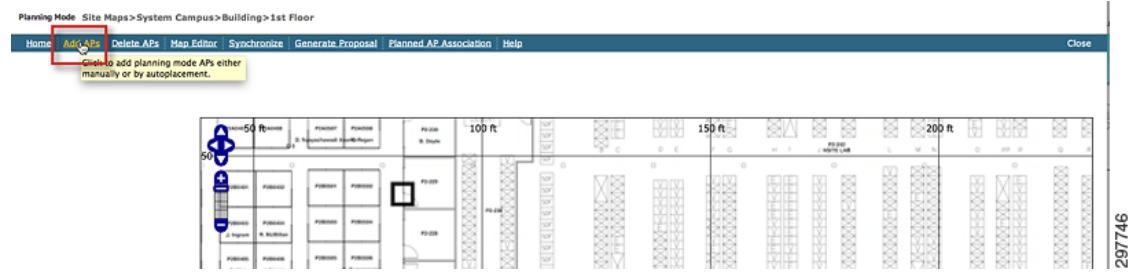


297745

- Step 5** In the new window that opens for RF Planner, click **Add APs** to begin positioning the APs.

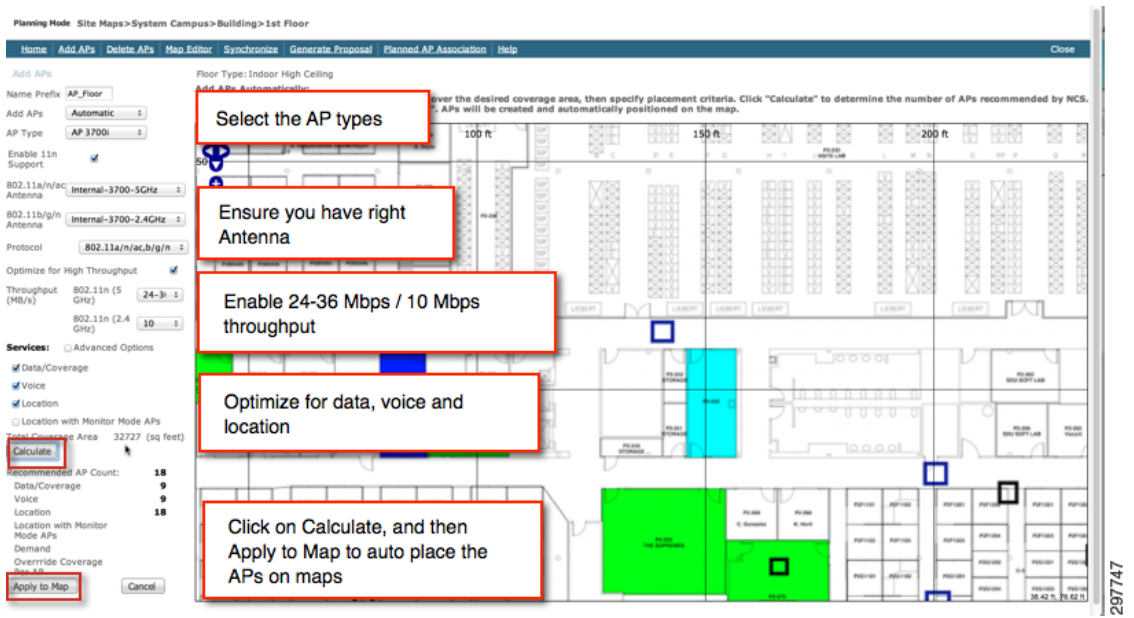


Figure 16-5 Position APs



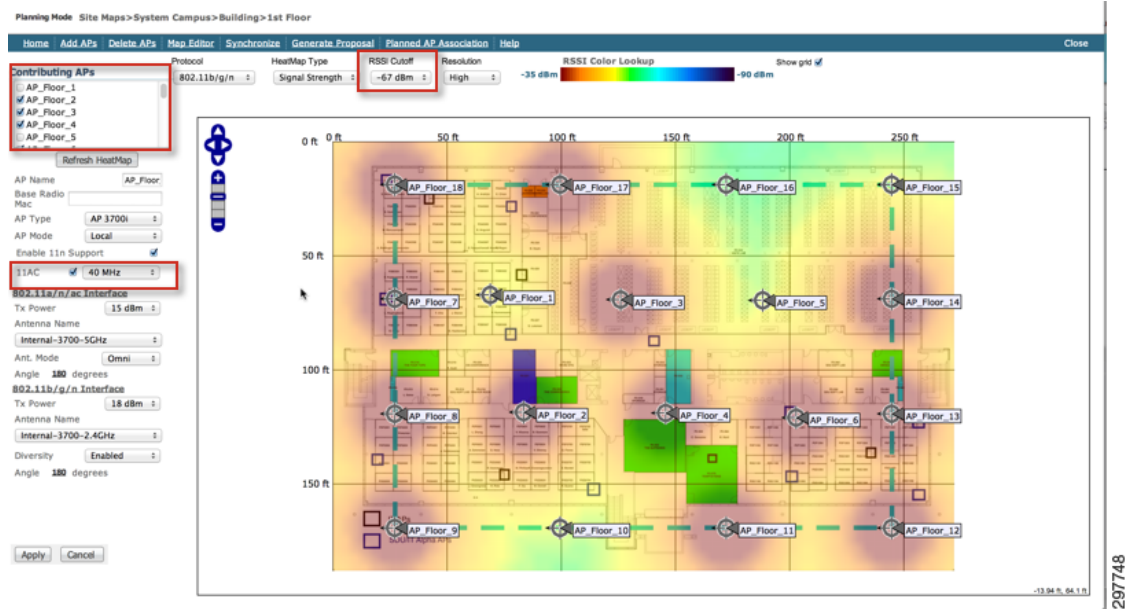
**Step 6** Follow the steps as outlined. Provide a name for the AP (optional), select the correct type of APs and the correct type of antennas, optimize for Data, Voice, and Location for CMX, and click **Calculate**. The Calculator then takes into account the various parameters that you input and suggests the number of APs that are required to obtain the desired coverage. Once satisfied, click **Apply on Map** to automatically place the APs on the maps.

Figure 16-6 Input AP Details



**Step 7** Once an initial RF Plan has been made, APs are put on the map automatically. Notice that the placement is optimized for both perimeter and indoor coverage.

Figure 16-7 APs Placed on Maps via RF Planner

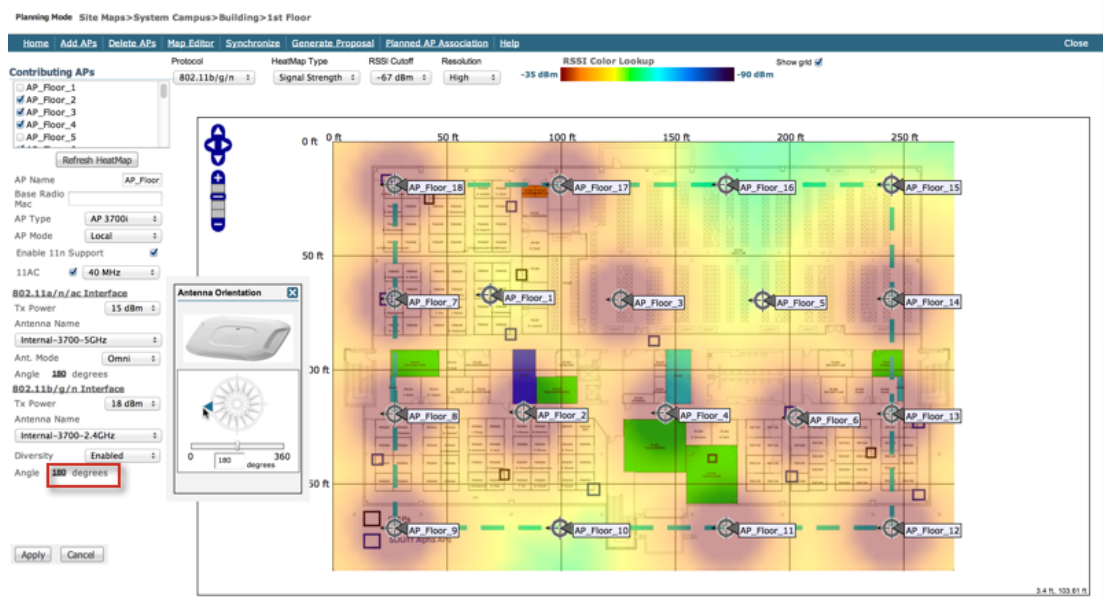


After the initial RF plan is completed, you can:

- Move APs around to meet your requirements.
- Delete APs by unchecking box on top left corner box for “Contributing APs”.
- If your deployment is an 802.11ac deployment, ensure that 802.11ac is selected.
- For RSSI threshold, the CMX solution recommends that the cut off be -67dbm to ensure that smartphones and tablets which have usually less powerful network cards than laptops are also clearly detected by the CMX solution.
- Click **Apply** to re-plan the RF plan.

**Step 8** Antenna azimuth angles for each AP can be adjusted by clicking the **Angle** link in the tool. This is useful if you have directional antennas.

Figure 16-8 AP Antenna Details



297749

**Step 9** Adjust the RF plan until you obtain a plan that meets your requirements.



**Note**

If you have walls and obstacles configured on the floor plan, then the RF plan tries to take that into account while planning. The type of floor (indoor high ceiling, office dry walls, etc.) also matters for RF planning. You may discover that you need fewer or more APs to get full coverage. Predictive RF Planner is a good tool to estimate and plan your network, but should not be used without a proper site survey.

**Step 10** Generate a RF Plan report if you want a copy of the plan. This report contains all the details about the recommended plan.

Figure 16-9 Generate Report



297750

# Ekahau RF Planning

The Ekahau Predictive Site Survey tool offers way to simulate RF planning, however it is **not** a replacement for a physical site survey as the tool cannot determine actual interferences and RF characteristics that may impact a WLAN. However the RF tool can be used to plan for deployment and Ekahau can be used for RF planning in deployments where a Cisco Prime Infrastructure is not available.


**Note**

To get started, download the latest version of the Ekahau Site Survey Planner from <http://www.ekahau.com/> (there is a fee).

The basic steps to use Ekahau for RF planning are:

- 
- Step 1** Click **Map > Add Map** button to import the floor map into the tool.
  - Step 2** Click the **scale** icon and set scale (to change the scale unit between meters/feet, go to **File > Preferences > Length Unit**).
  - Step 3** Set the Regulatory Domain (**File > Preferences > Regulatory Domain**).
  - Step 4** Click the **wall** icon to draw the walls (concrete on the outer boundary/stairs). This information should be available from background work you performed.
  - Step 5** Click the **simulated AP** icon and select the AP; for the purpose of this guide, we selected AP 3700.
  - Step 6** Click the **coverage area** icon and select boundaries on the map.
  - Step 7** Click on the **Auto Planner**. In the Coverage Requirement section, selection a location where CMX needs to be deployed. Next, use the values below as reference for various RF characteristics that can be configured in the tool:
    - Enter Signal Strength is -67 dBm
    - SNR ratio is 20 dBm
    - Data rate is 12
    - Number of Access Point = 3
    - Ping round Trip = 500
    - Packet Loss = 10%
  - Step 8** Click **Consider in Plan** under Capacity Requirements if you know the type and quantity of devices specified in the requirements. Click **Edit** and add that information.
  - Step 9** Click **Access Point type** and select the correct AP. Sometimes this goes reverts to a default, so be aware of this behavior.
  - Step 10** Click **Optimize Coverage for** and select either 2.4 or 5 Ghz.
  - Step 11** **Important:** Click **Advanced Settings** and change transmit power to 10mW .


**Note**

Regarding transmit power, even though APs can transmit maximum power up-to 18 mw (or 20mw in some cases), it is a good idea to base your RF plan on half power (i.e., 10 mw). This ensures that you are planning for location and mitigating effects of possible unknown and known interferences sources near the deployment. This also helps later when RRM is used to plan for channels.

- Step 12** You can change Antenna height or leave as default if not known.



**Step 13** If an 802.11ac AP is chosen, change **Bandwidth (on 5Ghz)** to 80 Mhz if required. This is not required for the CMX solution or in general for Wi-Fi design. Plan on using 80Mhz only if you are sure that all your clients also support 80Mhz.

**Step 14** Click **Create Plan**.

The APs now appear on the heatmap.

**Step 15** Click **View** and select **Access Point Names** to see the AP names on the map (recommended). You can unclick **Radio Channels** so they do not appear on the map.

The AP heat maps provide you with a rough estimate of how many APs might be needed for an installation. Note that this number may be less or more than what might be desired. Rearrange the APs to get a better idea of how many APs might be needed to satisfy the use case.

Calculate the area manually (e.g, 300 feet x 150 feet = 33,750 sq feet)

- Rule of thumb is 1 AP per:
  - 5000 sq feet for Data
  - 2500 sq feet for Voice/Location (CMX Deployment)
- So if we only do data then:  $33750/5000 = \text{need } 7 \text{ APs!}$

You can re-arrange AP location (remember to click **Arrow Pointer** first, so we do not add more APs).

You can also manually add APs if needed for better coverage.

**Step 16** If everything looks satisfactory with the heatmap, save the map:

Click **File > Export Image** (perhaps save in the same folder created for the case).

Click **File > Report** (you may leave all options selected for your reference).

If you have more maps to create, follow the steps above for each one. Remember that for each map, we need the Heatmap Image and the Report.

If there a multi-floor building, create a separate map for each floor.

For multiple AP Models, you need to manually click and select them. You cannot select two different AP Models on Auto planner.

---

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

- Perform a walk-around of the site and verify that areas on the floor plan where access point mounting is desired can actually accommodate it. This is always a good idea, since floor plans and blueprints do not always indicate the precise conditions present at each location where an access point may be mounted. For example, you may find that certain locations that appear to be viable candidates on paper actually are inaccessible (such as an electrical closet), inappropriate (such as an outdoor balcony), or are otherwise not acceptable. In such cases, access points should be relocated close to the original location such that the impact on the overall design is minimal.
- Verify RF propagation and coverage assumptions by temporarily installing a few access points in various test areas of the floor and measuring actual RF signal strength and cell-to-cell overlap using a portable client device with appropriate site survey software tools. This is an excellent time to measure the ambient noise levels of the potential access point cells as well and determine whether the projected signal to noise ratio will be sufficient. Note that Cisco's RRM feature also monitors client SNR and increases access point power if a number of clients are noticed to fall below a prescribed SNR threshold. For more information about RRM, refer to Radio Resource Management under Unified Wireless Networks at:

- [http://www.cisco.com/en/US/tech/tk722/tk809/technologies\\_tech\\_note09186a008072c759.shtml](http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml)
- <http://sdu-sharepoint/mobility/EM/Guest%20Access>
- Validate whether there are any radar users present in your locale that may interfere with the use of the additional 802.11a that are subject to DFS. If there are not, these channels can be made available for use by enabling DFS on your WLAN controllers.



## Multi-Floor Deployments

---

September 4, 2014

This chapter discusses challenges in deployments that involve multiple floors. Recommendations on what to keep in mind while designing for RF network are also discussed.

Not every venue is a simple office space with uniform ceilings and rectangular shapes. Furthermore, most venues have aesthetic restrictions that limit the flexibility in placing APs.

Many venues have high ceilings that overlook many floors, which is quite common in museums, theaters, malls, etc. where there is an atrium in the middle of the building. The MSE/Cisco Prime Infrastructure can handle high ceiling deployments quite easily with the proper RF profile selected. For every floor, the “Indoor High Ceiling” RF model can be specified in the Cisco Prime Infrastructure interface, as discussed in [Cisco Prime Infrastructure RF Planning Tool](#) in [Chapter 16](#), “[Predictive Radio Frequency Planning](#).”

Inter-floor issues arise mainly due to:

- Construction and structure of the building
- Placement of the APs

### Limited Flexibility for Placing APs

High ceiling deployments have challenges with AP placement. An open, high ceiling area might limit the options for placing APs. For example, the aforementioned ceiling mounted APs that is preferred may not be an option at all and APs mounted at lower height might be the only option.

The optimal position for an AP based on our measurements may be in an area that adversely impacts the aesthetics of the venue and the venue interior designers may not allow placement of an AP at the preferred location. It may also be impossible to run power and Ethernet wiring to the intended AP location.

### Inter-floor Interference Issues

In many situations, an atrium passes through multiple floors. In such venues, a wireless client on Floor-X can see APs from the floors above and below, leading to miscalculation of the floor level. This is what is usually labeled as Inter-floor interference. Basically the APs on multiple floors detect signals from a device and report their observed RSSIs to the MSE. The MSE uses these readings to calculate the

location and may end up placing the client on the wrong floor as the APs on the floor above or below the actual floor that the device belongs to may report RSSIs as strong as the APs on the actual floor. This leads to user confusion and frustration.

## AP Deployment Guidelines to Mitigate Inter-floor Issues

It is recommended to deploy vertically aligned APs to give best results for inter-floor sensitive location.

- APs can be one or two meters offset per floor but in general they are in alignment.
- Design location deployment independently for every floor, only taking lift shafts, stairwells, and atriums into consideration that impact RF bleed-over.
- Avoid placing APs directly in short range of any atrium where RF can be affected depending on reflections or attenuation. Where possible place APs so as to avoid line of site to adjacent floors or other levels.
- Having APs vertically aligned on adjacent floors is not a problem. In most modern deployments, dynamic channel and power contention can be resolved by the controllers. If co-channel interference or signal strength is greater than recommended levels, then monitor mode APs should be considered.
- Care should also be given to staggered AP deployments where floor materials have weak RF absorption.

Based on past experience, here are some of the ways to eliminate inter-floor interference issues. Before embarking on this tedious, iterative process, perform several walking tests with more than one device to be absolutely sure that inter-floor issues exist and that it must be resolved. Note that the steps below can be carried out locally around the area where inter-floor interference is experienced. The floor determination on adjacent floors should also be checked before and after implementing the following recommendations:

- Try to move APs—The first, easy step in mitigating inter-floor impact is to move APs, if possible, without impacting Wi-Fi and Location coverage. This is easier said than done due to challenges in AP placement and also the inability to move them at will. Moving APs 10-20 ft. away from atriums can have a large impact on location inter-floor issue and minimal impact on Wi-Fi coverage.
- Exclude APs from location calculation—This is a logical operation that does **not** disturb the physical APs or their positions or even the Wi-Fi connectivity in an area. In this approach, the floor maps are manipulated and the APs causing inter-floor issues are removed from the floor on the Cisco Prime Infrastructure. With the removal of the AP from the floor map, the MSE never takes into account any RSSI reading from that AP that is no longer present on the floor map, even when the WLC sends the RSSI reading to the MSE. This clever approach has a few benefits.
  - It does not impact Wi-Fi coverage.
  - It does not mandate any changes in the physical AP placement.
  - It does not have any impact on the AP or the WLC.
  - If it does not work, it can always be reversed easily with configuration changes from the Cisco Prime Infrastructure.

Obviously extensive testing must be performed across multiple floors when this approach is adopted to resolve inter-floor issues.

- Eliminate improbable location for clients (for example, atriums)—In many venues, there are areas that are unreachable for humans. A multi-level atrium is a classic example of an area that spans multiple floors but only a small part of the lowest level floor is accessible. The floor maps will still show the atrium as being part of every floor (see [Figure 17-1](#) and [Figure 17-2](#) of a museum). The

floor map for the Lower Level and Entry level look very similar with an area in the middle labeled “Great Hall”. This Great Hall area is an open area that is visible to all floors in the museum, but it is only accessible from the Lower Level floor.

Figure 17-1 Venue with Atrium


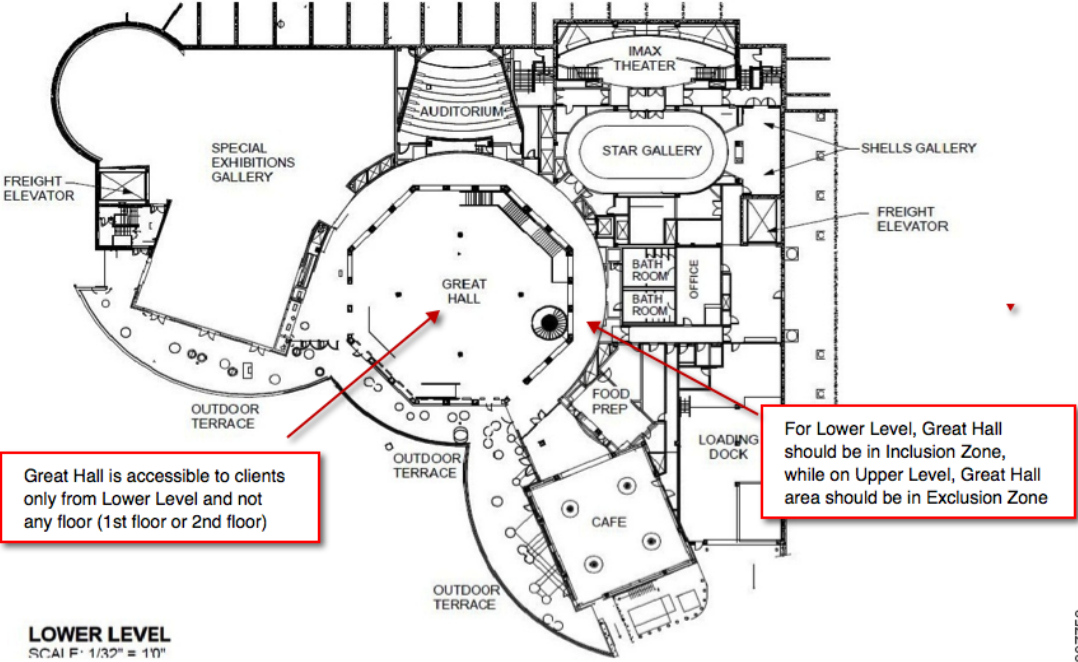


Figure 17-2 Lower Level and Upper Level Inclusion Zone

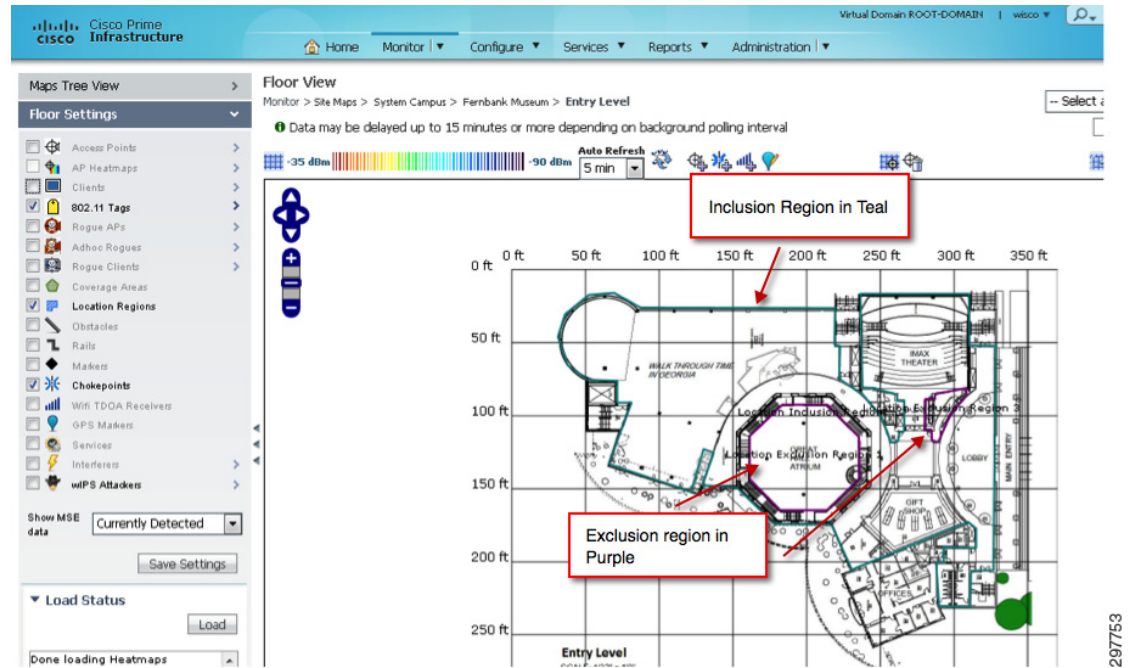


As can be seen from Figure 17-1 and Figure 17-2, a human can only be located within the “Great Hall” area only when present on the Lower Level. However due to the placement of the APs and Wi-Fi coverage, the MSE might calculate the location of a device to be on the Entry Level floor and within the Great Hall area. This is perfectly correct as per the MSE algorithms and setup, but not correct given the physical layout. Therefore the MSE must be configured to never locate a device within the Great Hall region when on the Entry Level (and Upper Level) floors. This is accomplished by drawing an Exclusion region on the floor map. By drawing such a region, we ensure that no device is ever located inside that area, even when the calculations fall within the region. This avoids annoyances and frustration for humans using an application with the MSE.

Figure 17-3 shows a picture of the museum floors with an exclusion region.



Figure 17-3 Inclusion and Exclusion Regions



**Note** More details on configuring inclusion and exclusion regions for a CMX deployment is covered in [Inclusion and Exclusion Areas on a Floor](#) in [Chapter 24, “Configuring Cisco Prime Infrastructure.”](#)

- Limiting where clients can be located—The opposite of the previous discussion on Exclusion regions are Inclusion regions. The idea is to restrict/confine the region in which an MSE can place clients. For example, as part of normal calculation of Location, if the MSE determines that a client is outside the building (which is not likely in most cases), then the MSE is forced to place the client within the Inclusion region. Continuing with the museum example, an Inclusion region has been drawn around the entire floor plan so that MSE always places the clients within the floor plan.

## Multi-Floor RF Site Survey

There are numerous factors that need to be taken into account when you survey multi-floor buildings, hospitals, and warehouses. It is important to find as much detail as possible in regard to the building construction. Some examples of typical construction methods and materials that affect the range and coverage area of APs include metallic film on window glass, leaded glass, steel-studded walls, cement floors and walls with steel reinforcement, foil-backed insulation, stairwells and elevator shafts, plumbing pipes and fixtures, and many others. Also, various types and levels of inventory can affect RF range, particularly those with high steel or water content. Some items to watch for include printer paper, cardboard boxes, pet food, paint, petroleum products, engine parts, and so forth. Ensure that the site survey is conducted at peak inventory levels or at times of highest activity. A warehouse at a 50% stocking level displays a very different RF footprint than the same facility that is completely occupied.

Similarly, an office area that is not populated has a different RF footprint than the same area when occupied. Although many parts of the site survey can be conducted without full occupation, it is essential to conduct the site survey verification and tweak key values at a time when people are present and normal activity takes place.

The higher the utilization requirements and the higher the density of users, the more important it is to have a well-designed diversity solution. When more users are present, more signals are received on the device of each user. Additional signals cause more contention, more null points, and more multipath distortion. Antenna diversity on the AP helps to minimize these conditions.

Keep these guidelines in mind when you conduct a site survey for a typical multi-floor building:

- Elevator shafts block and reflect RF signals.
- Supply rooms with inventory absorb RF signals.
- Interior offices with hard walls absorb RF signals.
- Kitchens can produce 2.4GHz interference caused by microwave ovens.
- Test labs can produce 2.4 GHz or 5 GHz interference. The problem of interference is that it increases the noise floor and decreases the SNR (signal to noise ratio) of the received signal. A higher noise floor reduces the effective range of the APs.
- Office cubicles tend to absorb and block signals.
- Class windows and partitions reflect and block RF signals.
- Bathroom tiles can absorb and block RF signals.
- Conference rooms require high AP coverage because they are a high Wi-Fi utilization area.

When you survey multi-floor facilities, APs on different floors can interfere with each other as easily as APs located on the same floor. This can be beneficial for voice and/or data deployments, but it causes problems when you deploy Context Aware. Floor separation is critical for this solution to function properly. In multi-tenant buildings, there can be security concerns that require the use of lower transmission powers and lower gain antennas to keep signals out of nearby rooms or offices.

## Hospitals

The survey process for a hospital is much the same as that for an enterprise, but the layout of a hospital facility tends to differ in these ways:

- Hospital buildings often have recurrent reconstruction projects and additions. Each additional construction can require different construction materials with different levels of signal attenuation.
- Signal penetration through walls and floors in the patient areas is typically minimal, which helps create micro-cells. Consequently, AP density needs to be much higher to provide sufficient RF coverage.
- The need for bandwidth increases with the increased usage of WLAN ultrasound equipment and other portable imaging applications.
- Due to the requirement for higher AP density, cell overlap can be high, which results in channel reuse.
- Hospitals can have several types of wireless networks installed, which includes 2.4 GHz non-802.11 equipment. This equipment can cause contention with other 2.4 GHz or 5 GHz networks.
- Wall-mounted diversity patch antennas and ceiling-mounted diversity omni-directional antennas are popular, but keep in mind that diversity is required.



## Warehouses

Warehouses have large open areas that often contain high storage racks. Many times these racks reach almost to the ceiling where APs are typically placed. Such storage racks can limit the area that the AP can cover. In these cases, consider placing APs on other locations besides the ceiling, such as side walls and cement pillars.

Also consider these factors when you survey a warehouse:

- Inventory levels affect the number of APs needed. Test coverage with two or three APs in estimated placement locations.
- Unexpected cell overlaps are likely because of coverage variations. The quality of the signal varies more than the strength of that signal. Clients can associate and operate better with APs farther away than with nearby APs.
- During a survey, APs and antennas usually do not have an antenna cable that connects them, but in a production environment, the AP and antenna can require antenna cables. All antenna cables have signal loss. The most accurate survey includes the type of antenna to be installed and the length of cable to be installed. A good tool to use to simulate the cable and its loss is an attenuator in a survey kit.

## Manufacturing Facility

When you survey a manufacturing facility, it is similar to the surveillance of a warehouse. One key difference is that the ambient RF environment is much noisier in a manufacturing facility because of many more sources of RF interference. Also applications in a manufacturing facility typically require more bandwidth than applications used in a warehouse environment. These applications can include video imaging and wireless voice. Multipath distortion is likely to be the greatest performance problem in a manufacturing facility.

It is important that the site survey not only measures signal levels, but also generates packets and then reports packet errors in order to properly characterize the RF environment.

For areas where user traffic is high, such as office spaces, schools, retail stores, and hospitals, Cisco recommends that you place the AP out of sight and place unobtrusive antennas below the ceiling.

**Note**

For more details, see:

<https://supportforums.cisco.com/sites/default/files/legacy/1/7/2/41271-Cisco%20Mobility%20Services%20Engine%20-Context%20Aware%20Mobility%20Solution%20Deployment%20Guide.pdf>





# Capacity Planning and High Density

September 4, 2014

This chapter discusses planning a network while keeping capacity and application requirements in mind. Today’s WLAN needs are heavily dependent on mobile devices and applications. Capacity planning involves looking at application needs and designing a network around them, while High Density networks may be required where far too many clients are expected to connect in a location.

Typically the most frequent question that a network designer faces is, “How many Access Points do I need?” While RF planning tools like Ekahau may attempt to answer that question theoretically, and one can assume that the recommendation of the tools are within range of 15-20% error rate, it is still a very useful exercise to plan for capacity planning before rather than after deployment.

## Access Point Density

The RF network can support these applications based on the expected accuracy (delta between estimated and actual location) and currency (time between location estimates) required. The accuracy is a direct function of the AP density (and AP height) while the currency is a function of the AP density and client type. The types of applications generally supported given an approximate AP density are shown in [Table 18-1](#).

**Table 18-1** CMX Venue and Density Required

Application	Venue types	Density (sq. ft/AP)	Accuracy	Currency
Presence	Mall, airport, etc	10+K	9-18m	N/A
Proximity	Retail, etc	<2.5K	5-8m	~30s
Asset-tracking	Enterprise, mall, etc.	5K	7-12m	>> 1min
Mobility-tracking	Mall, airport, etc.	<2.5K	5-8m	~30s

Consider an area of 500 feet x 500 feet floor of a mall space = 250,000 sq. ft.

From above, for a mobility tracking a density of 2.5K square feet per AP can be assumed for location accuracy of 5-8m and location currency of ~30 seconds.

So total number of APs required, roughly, would be = 250,000/2500 = 100 Access Points.

Note that if the only requirement were to track presence of a client in a CMX, but not engage the client in any application activity, then the AP requirements would change to:

Presence requirements = 10k / AP with accuracy of 9-18m

Mall area = 250,000 qft

Number of APs required = 250,000/10000 = 25

It is recommended that for CMX deployment, the WLAN network be designed with mobility tracking in mind with data and voice as an additional requirement. The CMX solution not only enables presence analytics, but also provides ways of providing guest access and engaging the user. Hence it is highly recommended that designers plan for voice, data, and location instead of presence. While 2,500 sq. ft./ AP is generally considered a good recommendation that covers a wide range of use cases, it is by no means a fit for every situation. Different RF considerations, type of materials, and applications provided need to be considered before planning for a deployment.

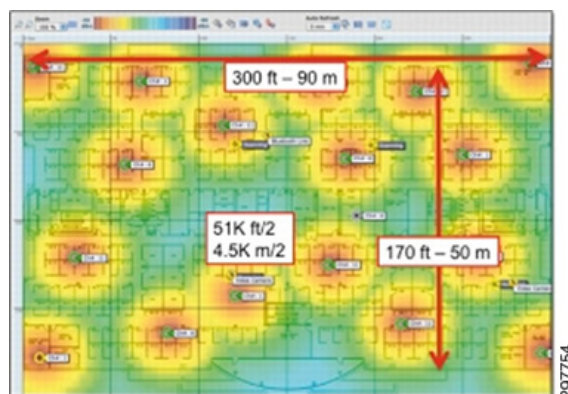
## High Density Deployment

While a detailed High Density Deployment discussion is beyond the scope of this validated design guide, here are a few guidelines when preparing for a high density deployment. A high density CMX deployment is a network that may have a higher number of Wi-Fi clients than a traditional cell may have. High-density is defined as any environment with a large concentration of users, such as a classroom, lecture hall, or auditorium where the users are connected wirelessly, sharing applications, and using other network services individually.

High-density WLAN design refers to any environment where client devices will be positioned in densities greater than coverage expectations of a normal enterprise deployment, in this case a traditional, carpeted office. For reference, a typical office environment has indoor propagation characteristics for signal attenuation. User density is the critical factor in the design. Aggregate available bandwidth is delivered per radio cell and the number of users and their connection characteristics (such as speed, duty cycle, radio type, band, signal, and SNR) occupying that cell determines the overall bandwidth available per user.

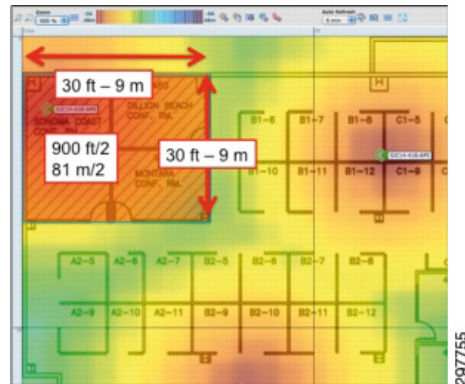
A typical office environment, as shown in [Figure 18-1](#), may have APs deployed for 2,500 to 5,000 square feet with a signal of -67 decibels in milliwatts (dBm) coverage and a maximum of 20 to 30 users per cell. That is a density of one user every 120 square foot (sq. ft.) and yields a minimum signal of -67 dBm at the edge of the cell.

**Figure 18-1** Typical Office WLAN



In planning and deploying such a WLAN, an AP is typically placed in an area expected to have a higher user density, such as in a conference room, while common areas are left with less coverage. In this way, pre-planning for high-density areas is anticipated. Conference rooms are often placed in clusters, so it is best to design for the maximum capacity of the area. For example, if maximum occupancy for the three rooms is 32, user density would be one user per 28 square feet, as shown in Figure 18-2.

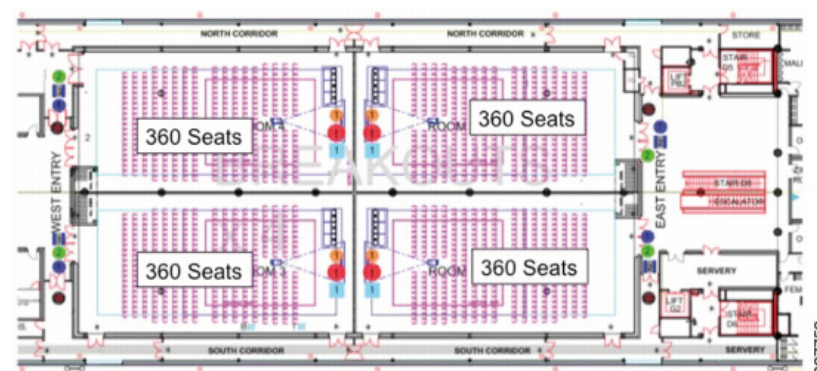
**Figure 18-2 Calculate User Density**



In a high-density environment such as a lecture hall or auditorium, the densities of users in the occupied space increase dramatically. User seating is typically clustered very close together to achieve high occupancy. The overall dimensions of the space are really only useful for getting an idea of the free space path loss of the AP signal. User densities are not evenly distributed over the entire space as aisle ways, stages, and podiums represent a percentage of space which is relatively unoccupied. The RF dynamics of the AP are very different from those experienced at the user level. The APs are exposed with an excellent view of the room and the user devices will be packed closely together with attenuating bodies surrounding them.

The single biggest sources of interference in the room are the client devices themselves. For each user sitting in the auditorium who can rest their hand comfortably on the back of the seat in front of them, the distance is approximately three feet, with an average seat width of 24 inches. This yields what is defined as a high-density environment, with less than 1 square meter per device deployed, assuming one or more devices connected per seat.

**Figure 18-3 Seating and Interference**



What is ultimately going to affect the client devices more than any other factor is the degradation of signal-to-noise ratio (SNR) through both co-channel and adjacent channel interference driven by co-located devices. Proper system engineering can minimize the impact by maximizing proper spatial

reuse, but it cannot be eliminated in highly dense environments entirely. Operating margins become more critical as space is condensed and a bad radio or behavior in the mix can have a large impact within a cell. Client behavior under these conditions will vary widely and trends based on environment and event type have also been reported. There is not much that can be done about the particular client mix or behavior. The design goal is to engineer the network side as robustly as possible and to control and understand all variables.

Within environments that qualify as high-density, there are also submodels built by use case. For example, in a high-density environment such as a public venue or stadium, capacity is planned based on what percentage of users are likely to be active on the network at any one time. In higher education there is a different model, where casual WLAN activity is one use case while activity when a professor is lecturing may increase dramatically, up to 100 percent.

The WLAN design process can begin in many ways, but generally it begins with an expressed desire to provide connections to a specific area where a number of users will participate in a focused activity. To evaluate what is possible, it is first necessary to understand what is required as well as what is possible. There is generally a primary application that is driving the need for connectivity. Understanding the throughput requirements for this application and for other activities that will take place on the network will provide the designer with a per-user bandwidth goal. Multiplying this number by the number of expected connections yields the aggregate bandwidth that will be required.

The required per connection bandwidth is used to drive subsequent design decisions.

## Establish and Validate a Per-Connection Bandwidth Requirement

How much bandwidth does each user require on average? In [Table 18-2](#), the nominal throughput requirements for several popular applications and use cases in a higher education setting are shown.

**Table 18-2** Bandwidth Requirements per Application

Application by Use Case	Nominal Throughput
Web—Casual	500 kilobits per second (Kbps)
Web—Instructional	1 Megabit per second (Mbps)
Audio—Casual	100 Kbps
Audio—Instructional	1 Mbps
On-demand or Streaming Video—Casual	1 Mbps
On-demand or Streaming Video—Instructional	2-4 Mbps
Printing	1 Mbps
File Sharing—Casual	1 Mbps
File Sharing—Instructional	2-8 Mbps
Online Testing	2-4 Mbps
Device Backups	10-50 Mbps

In all cases, it is highly advisable to test the target application and validate its actual bandwidth requirements. Software designers are often required to pick just one average number to represent the application's requirements when there are actually many modes and deployment decisions that can make

up a more accurate number. It is also important to validate applications on a representative sample of the devices that are to be supported in the WLAN. Additionally, not all browsers and operating systems enjoy the same efficiencies, and an application that runs fine in 100 kilobits per second (Kbps) on a Windows laptop with Microsoft Internet Explorer or Firefox, may require more bandwidth when being viewed on a smart phone or tablet with an embedded browser and operating system.

Once the required bandwidth throughput per connection and application is known, this number can be used to determine the aggregate bandwidth required in the WLAN coverage area. To arrive at this number, multiply the minimum acceptable bandwidth by the number of connections expected in the WLAN coverage area. This yields the target bandwidth needed for the need series of steps.

## Calculate the Aggregate Throughput Required for the Coverage Area

In a WLAN, a channel's speed is affected by multiple factors including protocols, environmental conditions, and operating band of the adapter. Before calculating aggregate throughput, several things must be considered.

In the aggregate throughput calculation, the connections instead of the seats were used as the basis for calculation. The number of connections in a cell is what determines the total throughput that will be realized per connection instead of the number of seats. Most users today carry both a primary computing device (such as a smartphone, tablet computer, or laptop) as well as a second device (such as a smartphone). Each connection operating in the high-density WLAN consumes air time and network resources and will therefore be part of the aggregate bandwidth calculation. An increase in numbers of device connections is one of the primary reasons older WLAN designs are reaching oversubscription today.

Users and applications also tend to be bursty (a measure of the unevenness or variations in the traffic flow) in nature and often access layer networks are designed with a 20:1 oversubscription to account for these variances. Application and end user anticipated usage patterns must be determined and also accounted for. Some applications, such as streaming multicast video, drive this oversubscription ratio down while others may drive this factor even higher to determine an acceptable SLA for each cell's designed capacity.

For 802.11 wireless networks or any radio network in general, air is the medium of propagation. While there have been many advances in efficiency, it is not possible to logically limit the physical broadcast and collision domain of an RF signal or separate its spectrum footprint from other radios operating in the same spectrum. For that reason, Wi-Fi uses a band plan that breaks up the available spectrums into a group of non-overlapping channels. A channel represents a cell. Using the analogy of Ethernet, a cell represents a single contiguous collision domain.

How many users can access an AP comfortably? Hundreds. But the question should not be how many users can successfully associate to an AP but how many users can be packed into a room and still obtain per-user bandwidth throughput that is acceptable. The question revolves back to around what is the expectation of a CMX deployment.

## 802.11 and Scalability—How Much Bandwidth Will a Cell Provide?

To scale 802.11 networks to reliably deliver consistent bandwidth to a large number of users in close proximity, it is important to examine certain WLAN fundamentals under reasonably ideal conditions. Once the rules are understood, the ways to manipulate them to maximum advantage are presented.

In real WLANs, the actual application throughput is what matters to the end user and this differs from the signaling speed. Data rates represent the rate at which data packets are carried over the medium. Packets contain a certain amount of overhead that is required to address and control the packets. The application throughput is carried as payload data within that overhead. [Table 18-3](#) shows average application throughput by protocol under good RF conditions. Account for and manage all energy within the operating spectrum to ensure all of it is available for use

**Table 18-3** Average Application Throughput by Protocol

Protocol	Throughput (Mbps)
802.11b	7.2
802.11b/g mix	13
802.11g	25
802.11a	25
802.11n—HT20 one spatial stream (1ss) Modulation Coding Scheme 7 (MCS7)	25
802.11n—HT20 2ss Modulation Coding Scheme 15 (MCS15)	70
802.11n—HT40 2ss Modulation Coding Scheme 15 (MCS15)	160
802.11ac—HT80, 1ss Modulation Coding Scheme 8, Short Guard Interval (SGI)	390
802.11ac—HTC80, 2ss Modulation Coding Scheme 8, Short Guard Interval (SGI)	867
802.11ac—HTC80, 3ss Modulation Coding Scheme 8, Short Guard Interval (SGI)	1.3 Gbps

The discussion until now has centered on a use case where every client in the room will be competing for bandwidth simultaneously. This is the case when the users in the room simultaneously access a resource on queue. However there are many instances where the design requirement is to offer access to resources or the Internet for casual use at an event or within a venue such as a sports arena. Planning and sizing for these types of events can be quite different and is based on expected Client Duty Cycle.

At a sporting event, for example, there are certain areas that require ubiquitous and instant access during the entire event. Ticketing, vendor sales, staff, and press areas generally require the highest amount of access. Of these, the press area is the only one that requires a high level of capacity in the arena itself. For the fans attending the event, only a percentage will be active on the WLAN at any one time. From experience we see a 20 to 30 percent take rate with some well-defined peaks occurring during period breaks. During play, very few fans are accessing the WLAN. However this is changing as applications such as video replay, instant stats, and concession orders from the seat become more commonplace.

Observation and understanding of the requirements of WLAN users and situational requirements guide the development of reasonable design goals. 500 users in a room who require simultaneous access to a single resource is a different design challenge than 1,000 or 1,500 users who only occasionally use the wireless network. Also, be aware that user patterns can and do change with time. This has been seen with the increase in the number of network clients per user. Monitoring network access and keeping good



statistics allows wireless engineers to stay on top of user trends on the university campus. Good management platforms such as Cisco Prime Infrastructure are essential for managing the resulting network in real time and monitoring trends in a proactive manner.

## Other High Density Considerations

- Choose a high minimum data rate to support increased efficiency, lower duty cycle, and reduce the effective size of the resulting cell.
- Plan for Wi-Fi co-channel interference.
- 5 GHz support is critical for high-density, so determine the channel plan that you will support and how it will be administered.
- Determine the number of channels and cells needed.



### Note

---

While an extensive discussion of high density deployment is beyond the scope of this validated design guide, designers interested in deploying the CMX solution in a high-density environment are highly encouraged to read the Cisco High Density Design Guide for additional design recommendations: [http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1250-series/design\\_guide\\_c07-693245.html#wp9001186](http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1250-series/design_guide_c07-693245.html#wp9001186).

---





## Location Voice and Data Co-Existence

---

September 4, 2014

This chapter discusses the pertinent characteristics of voice and data designs only as they relate to co-existence with the location tracking capabilities of the Cisco UWN.

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

When architecting VoWLAN and high speed data designs and determining subsequent access point placement, the primary concerns that the designer should consider are described below.

### Minimum Desired Cell Signal Level Threshold

For example, when designing VoWLAN solutions that involve the Cisco VoWLAN handset, current VoWLAN best practices suggest a minimum planned signal level threshold of -67dBm. Other voice devices may have differing requirements. Requirements for data devices depend on the transmission rate at which they are required to operate. Data devices used to pass streaming multimedia and other bandwidth-intensive applications typically require higher data transmission rates and consequently, higher minimum signal levels. This is covered as part of AP capacity planning (see [Chapter 18, “Capacity Planning and High Density”](#)).

### Signal to Noise Ratio (SNR)

This is the ratio of the signal strength at the receiver to the noise floor and is measured in dB. Since both components of the ratio are specified in dBm, the SNR can be calculated by simply subtracting the noise value from the signal strength value. The minimum required SNR for a receiver to operate properly varies depending on construction of the receiver as well as the bit rate or modulation it is expected to operate at. A typical example is shown in [Table 19-1](#).

**Table 19-1 Data Rate to SNR Ratio**

<b>Data Rate</b>	<b>2.4 GHz</b>	<b>5 Ghz</b>		
	Min RSSI	Min SNR	Min RSSI	Min SNR
14.4/30	-82	11	-79	14
28.9/60	-79	14	-76	17
43.4/90	-77	16	-74	19
57.8/120	-74	19	-71	22
86.7/180	-70	23	-67	26
115.6/240	-66	27	-63	30
130/270	-65	28	-62	31
144.4/300	-64	29	-61	32

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

## Data Rate

Data bit rates are enabled or disabled via the wireless infrastructure, with minimum signal level thresholds and the signal to noise ratio determining which of the enabled bit rates are actually usable. For example, with the Cisco VoWLAN handset, the combination of a -67dBm minimum signal level and a 25dB signal to noise ratio generally makes the use of 24 Mbps or greater data rates possible.

## Cell-to-Cell Overlap


In a very simple sense, we can think of each of our access points as residing at the center of an RF “cell” with a spherical boundary of RF coverage around them. Our primary interest is in the coverage boundary associated with our desired minimal signal threshold. To provide consistent coverage and availability across our floor, each of our cells should join with each adjacent cell at a coverage boundary that is greater than our desired minimal signal threshold. How much greater? That is determined by the amount of cell-to-cell overlap we wish to implement in our design, which in conjunction with the other parameters we have described, dictates the potential packet loss experienced by VoWLAN devices before a roam event occurs.

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

that devices have the opportunity to roam well before the quality of the user's voice call is in jeopardy. We accomplish this by ensuring that the recommended degree of cell-to-cell overlap exists in our designs.


Figure 19-1 illustrates the concept of cell overlap for a smartphone or tablet using 802.11bgn. For CMX, the recommended best practices suggest that the cell-to-cell overlap should be approximately 20 percent when using 802.11bgn and approximately 15 percent when using 802.11an/ac.

**Figure 19-1 20% Inter Cell Overlap**



Data applications, on the other hand, typically do not display the same level of sensitivity to packet loss as do voice applications, hence they seldom require the same degree of cell-to-cell overlap. In most cases, a minimum 10% cell-to-cell overlap is sufficient for reliable roaming with data applications, as illustrated in Figure 19-2. High speed data applications and applications combining voice and data capabilities in a single device (smartphones, for example) may require cell-to-cell overlap that resembles a VoWLAN design much more than a data design.

**Figure 19-2 10% Inter Cell Overlap**



As seen, the designer striving towards completing an optimized location design is likely to be attempting to satisfy the four primary concerns of VoWLAN and data WLAN designers concurrently.

The chief location-tracking concerns of most designers wishing to track clients centers around:

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

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

The question that comes to mind then is can the requirements described earlier for voice and data applications be met in combination with the requirements of location tracking? The answer is yes, with the precise mechanics of how it is done dependent upon the specific requirements of the voice and data applications themselves, the access point and antenna configuration being considered, and the physical characteristics of the environment into which the infrastructure will be deployed.

**Note**

---

An example use case of a layout that needs to be updated from a VoWLAN ready to CMX (that is Location, Data, and Voice ready) is described in [Appendix D, “CMX Use Case Example—Upgrade VoWLAN Ready Network to Location/CMX Ready.”](#)

---



## Post-Deployment Radio Frequency Tuning

---

September 4, 2014

This chapter discusses post deployment RF tuning that should be done regularly on the deployment and includes using RRM for channel planning, CleanAir to mitigate RF interference, and a regular post-site survey assessment to ensure that optimum RF health is maintained.

Once a CMX solution is deployed, it is highly recommended that administrators turn on both Radio Resource Management and CleanAir to plan, optimize, and mitigate interference in the network.

### Radio Resource Management

The Radio Resource Management (RRM) software embedded in the Cisco Wireless LAN Controller acts as a built-in RF engineer to consistently provide real-time RF management of your wireless network. RRM enables Cisco WLCs to continually monitor their associated lightweight access points for the following information:

- Traffic load—The total bandwidth used for transmitting and receiving traffic, which enables wireless LAN managers to track and plan network growth ahead of client demand.
- Interference—The amount of traffic coming from other 802.11 sources.
- Noise—The amount of non-802.11 traffic that is interfering with the currently assigned channel.
- Coverage—The received signal strength (RSSI) and signal-to-noise ratio (SNR) for all connected clients.
- Other—The number of nearby access points.

Using this information, RRM can periodically reconfigure the 802.11 RF network for best efficiency. To do this, RRM performs these functions:

- Radio resource monitoring
- Transmit power control
- Dynamic channel assignment
- Coverage hole detection and correction

RRM automatically detects and configures new Cisco WLCs and lightweight access points as they are added to the network. It then automatically adjusts associated and nearby lightweight access points to optimize coverage and capacity.

Lightweight access points can simultaneously scan all valid 802.11a/b/g/n/ac channels for the country of operation as well as for channels available in other locations. The access points go “off-channel” for a period not greater than 60 ms to monitor these channels for noise and interference. Packets collected during this time are analyzed to detect rogue access points, rogue clients, ad-hoc clients, and interfering access points.

**Note**

In the presence of voice traffic (in the last 100 ms), the access points defer off-channel measurements.

Each access point spends only 0.2 percent of its time off-channel. This activity is distributed across all access points so that adjacent access points are not scanning at the same time, which could adversely affect wireless LAN performance.

**Note**

When there are numerous rogue access points in the network, the chance of detecting rogues on channels 157 or 161 by a FlexConnect or local mode access point is small. In such cases, the monitor mode AP can be used for rogue detection.

## Transmit Power Control

The Cisco WLC dynamically controls access point transmit power based on real-time wireless LAN conditions. You can choose between two versions of transmit power control: TPCv1 and TPCv2. With TPCv1, typically power can be kept low to gain extra capacity and reduce interference. With TPCv2, transmit power is dynamically adjusted with the goal of minimum interference. TPCv2 is suitable for dense networks. In this mode, there could be higher roaming delays and coverage hole incidents.

The Transmit Power Control (TPC) algorithm both increases and decreases an access point’s power in response to changes in the RF environment. In most instances, TPC seeks to lower an access point’s power to reduce interference, but in the case of a sudden change in the RF coverage—for example, if an access point fails or becomes disabled—TPC can also increase power on surrounding access points. This feature is different from coverage hole detection, which is primarily concerned with clients. TPC provides enough RF power to achieve desired coverage levels while avoiding channel interference between access points.

Table 20-1 shows the power level mapping.

**Table 20-1 Power Level Table**

Power Level	2.4 GHz	5 GHz
1	23 dBm (200 mW) CCK Only	20 dBm (100 mW)
2	20 dBm (100 mW)	17 dBm (50 mW)
3	17 dBm (50 mW)	14 dBm (25 mW)
4	14 dBm (25 mW)	11 dBm (12.5 mW)
5	11 dBm (12.5 mW)	8 dBm (6.25 mW)
6	8 dBm (6.25 mW)	5 dBm (3.13 mW)
7	5 dBm (3.13 mW)	2 dBm (1.56 mW)
8	2 dBm (1.56 mW)	-1 dBm (0.78 mW)
	-1 dBm (0.78 mW)	



# Overriding the TPC Algorithm with Minimum and Maximum Transmit Power Settings

The TPC algorithm balances RF power in many diverse RF environments. However it is possible that automatic power control will not be able to resolve some scenarios in which an adequate RF design was not possible to implement due to architectural restrictions or site restrictions—for example, when all access points must be mounted in a central hallway, placing the access points close together, but requiring coverage out to the edge of the building.

In these scenarios, you can configure maximum and minimum transmit power limits to override TPC recommendations. The maximum and minimum TPC power settings apply to all access points through RF profiles in a RF network.

To set the Maximum Power Level Assignment and Minimum Power Level Assignment, enter the maximum and minimum transmit power used by RRM in the text boxes in the Tx Power Control page. The range for these parameters is -10 to 30 dBm. The minimum value cannot be greater than the maximum value; the maximum value cannot be less than the minimum value.

If you configure a maximum transmit power, RRM does not allow any access point attached to the controller to exceed this transmit power level (whether the power is set by RRM TPC or by coverage hole detection). For example, if you configure a maximum transmit power of 11 dBm, then no access point would transmit above 11 dBm, unless the access point is configured manually.

## Dynamic Channel Assignment

Two adjacent access points on the same channel can cause either signal contention or signal collision. In a collision, data is not received by the access point. This functionality can become a problem, for example, when someone reading e-mail in a cafe affects the performance of the access point in a neighboring business. Even though these are completely separate networks, someone sending traffic to the cafe on channel 1 can disrupt communication in an enterprise using the same channel. Controllers can dynamically allocate access point channel assignments to avoid conflict and to increase capacity and performance. Channels are “reused” to avoid wasting scarce RF resources. In other words, channel 1 is allocated to a different access point far from the cafe, which is more effective than not using channel 1 altogether.

The controller’s Dynamic Channel Assignment (DCA) capabilities are also useful in minimizing adjacent channel interference between access points. For example, two overlapping channels in the 802.11b/g band, such as 1 and 2, cannot both simultaneously use 11/54 Mbps. By effectively reassigning channels, the controller keeps adjacent channels separated.

**Note**

---

We recommend that you use only non-overlapping channels (1, 6, 11, and so on).

---

The controller examines a variety of real-time RF characteristics to efficiently handle channel assignments as follows:

- Access point received energy—The received signal strength measured between each access point and its nearby neighboring access points. Channels are optimized for the highest network capacity.
- Noise—Noise can limit signal quality at the client and access point. An increase in noise reduces the effective cell size and degrades user experience. By optimizing channels to avoid noise sources, the controller can optimize coverage while maintaining system capacity. If a channel is unusable due to excessive noise, that channel can be avoided.

- **802.11 Interference**—Interference is any 802.11 traffic that is not part of your wireless LAN, including rogue access points and neighboring wireless networks. Lightweight access points constantly scan all channels looking for sources of interference. If the amount of 802.11 interference exceeds a predefined configurable threshold (the default is 10 percent), the access point sends an alert to the controller. Using the RRM algorithms, the controller may then dynamically rearrange channel assignments to increase system performance in the presence of the interference. Such an adjustment could result in adjacent lightweight access points being on the same channel, but this setup is preferable to having the access points remain on a channel that is unusable due to an interfering foreign access point. In addition, if other wireless networks are present, the controller shifts the usage of channels to complement the other networks. For example, if one network is on channel 6, an adjacent wireless LAN is assigned to channel 1 or 11. This arrangement increases the capacity of the network by limiting the sharing of frequencies. If a channel has virtually no capacity remaining, the controller may choose to avoid this channel. In very dense deployments in which all nonoverlapping channels are occupied, the controller does its best, but you must consider RF density when setting expectations.
- **Load and utilization**—When utilization monitoring is enabled, capacity calculations can consider that some access points are deployed in ways that carry more traffic than other access points (for example, a lobby versus an engineering area). The controller can then assign channels to improve the access point with the worst performance reported. The load is taken into account when changing the channel structure to minimize the impact on clients currently in the wireless LAN. This metric keeps track of every access point's transmitted and received packet counts to determine how busy the access points are. New clients avoid an overloaded access point and associate to a new access point. This parameter is disabled by default.

The controller combines this RF characteristic information with RRM algorithms to make system-wide decisions. Conflicting demands are resolved using soft-decision metrics that guarantee the best choice for minimizing network interference. The end result is optimal channel configuration in a three-dimensional space, where access points on the floor above and below play a major factor in an overall wireless LAN configuration.

**Note**


---

Radios using 40-MHz channels in the 2.4-GHz band or are not supported by DCA and cannot be configured.

---

The RRM startup mode is invoked in the following conditions:

- In a single-controller environment, the RRM startup mode is invoked after the controller is rebooted.
- In a multiple-controller environment, the RRM startup mode is invoked after an RF Group leader is elected.

You can trigger RRM startup mode from CLI.

RRM startup mode runs for 100 minutes (10 iterations at 10-minute intervals). The duration of the RRM startup mode is independent of the DCA interval, sensitivity, and network size. The startup mode consists of 10 DCA runs with high sensitivity (making channel changes easy and sensitive to the environment) to converge to a steady state channel plan. After the startup mode is finished, DCA continues to run at the specified interval and sensitivity.

## Coverage Hole Detection and Correction

The RRM coverage hole detection algorithm can detect areas of radio coverage in a wireless LAN that are below the level needed for robust radio performance. This feature can alert you to the need for an additional (or relocated) lightweight access point.

If clients on a lightweight access point are detected at threshold levels (RSSI, failed client count, percentage of failed packets, and number of failed packets) lower than those specified in the RRM configuration, the access point sends a “coverage hole” alert to the controller. The alert indicates the existence of an area where clients are continually experiencing poor signal coverage without having a viable access point to which to roam. The controller discriminates between coverage holes that can and cannot be corrected. For coverage holes that can be corrected, the controller mitigates the coverage hole by increasing the transmit power level for that specific access point. The controller does not mitigate coverage holes caused by clients that are unable to increase their transmit power or are statically set to a power level because increasing their downstream transmit power might increase interference in the network.

## Benefits of RRM

RRM produces a network with optimal capacity, performance, and reliability. It frees you from having to continually monitor the network for noise and interference problems, which can be transient and difficult to troubleshoot. RRM ensures that clients enjoy a seamless, trouble-free connection throughout the Cisco unified wireless network.

RRM uses separate monitoring and control for each deployed network: 802.11an/ac and 802.11bgn. The RRM algorithms run separately for each radio type (802.11an/ac and 802.11b/g). RRM uses both measurements and algorithms. RRM measurements can be adjusted using monitor intervals, but they cannot be disabled. RRM algorithms are enabled automatically but can be disabled by statically configuring channel and power assignment. The RRM algorithms run at a specified updated interval, which is 600 seconds by default.

For more detailed discussion and configuration of RRM, refer to:

[http://www.cisco.com/c/en/us/td/docs/wireless/controller/7-6/configuration-guide/b\\_cg76/b\\_cg76\\_chapter\\_01111110.html](http://www.cisco.com/c/en/us/td/docs/wireless/controller/7-6/configuration-guide/b_cg76/b_cg76_chapter_01111110.html).

## CleanAir

Cisco CleanAir is a spectrum intelligence solution designed to proactively manage the challenges of a shared wireless spectrum. It allows you to see all of the users of the shared spectrum (both native devices and foreign interferers). It also enables you or your network to act upon this information. For example, you could manually remove the interfering device or the system could automatically change the channel away from the interference. CleanAir provides spectrum management and RF visibility.

A Cisco CleanAir system consists of CleanAir-enabled access points, Cisco Wireless LAN Controllers, and Cisco Prime Infrastructure. These access points collect information about all devices that operate in the industrial, scientific, and medical (ISM) bands, identify and evaluate the information as a potential interference source, and forward it to the Cisco WLC. The Cisco WLC controls the access points, collects spectrum data, and forwards information to Cisco Prime Infrastructure or a Cisco mobility services engine (MSE) upon request.

For every device operating in the unlicensed band, Cisco CleanAir tells you what it is, where it is, how it is impacting your wireless network, and what actions you or your network should take. It simplifies RF so that you do not have to be an RF expert.

Wireless LAN systems operate in unlicensed 2.4- and 5-GHz ISM bands. Many devices, such as microwave ovens, cordless phones, and Bluetooth devices also operate in these bands and can negatively affect Wi-Fi operations.

Some of the most advanced WLAN services, such as voice over wireless and IEEE 802.11n radio communications, could be significantly impaired by the interference caused by other legal users of the ISM bands. The integration of Cisco CleanAir functionality into the Cisco Unified Wireless Network addresses this problem of radio frequency (RF) interference.

CleanAir is supported on mesh AP backhaul at a 5-GHz radio of mesh. You can enable CleanAir on backhaul radios and can provide report interference details and air quality.

## Role of the Cisco Wireless LAN Controller in a Cisco CleanAir System

The Cisco WLC performs the following tasks in a Cisco CleanAir system:

- Configures Cisco CleanAir capabilities on the access point.
- Provides interfaces (GUI, CLI, and SNMP) for configuring Cisco CleanAir features and retrieving data
- Displays spectrum data.
- Collects and processes air quality reports from the access point and stores them in the air quality database. The Air Quality Report (AQR) contains information about the total interference from all identified sources represented by the Air Quality Index (AQI) and summary for the most severe interference categories. The CleanAir system can also include unclassified interference information under per interference type reports, which enables you to take action in cases where the interference due to unclassified interfering devices is more.
- Collects and processes interference device reports (IDRs) from the access point and stores them in the interference device database.
- Forwards spectrum data to Prime Infrastructure and the MSE.

## Interference Types that Cisco CleanAir Can Detect

Cisco CleanAir can detect interference, report on the location and severity of the interference, and recommend different mitigation strategies. Two such mitigation strategies are persistent device avoidance and spectrum event-driven RRM.

Wi-Fi chip-based RF management systems share these characteristics:

- Any RF energy that cannot be identified as a Wi-Fi signal is reported as noise.
- Noise measurements that are used to assign a channel plan tend to be averaged over a period of time to avoid instability or rapid changes that can be disruptive to certain client devices.
- Averaging measurements reduces the resolution of the measurement. As such, a signal that disrupts clients might not look like it needs to be mitigated after averaging.
- All RF management systems available today are reactive in nature.

Cisco CleanAir is different and can positively identify not only the source of the noise but also its location and potential impact to a WLAN. Having this information allows you to consider the noise within the context of the network and make intelligent and, where possible, proactive decisions. For CleanAir, two types of interference events are common:

- Persistent interference
- Spontaneous interference

Persistent interference events are created by devices that are stationary in nature and have intermittent but largely repeatable patterns of interference. For example, consider the case of a microwave oven located in a break room. Such a device might be active for only 1 or 2 minutes at a time. When operating, however, it can be disruptive to the performance of the wireless network and associated clients. Using Cisco CleanAir, you can positively identify the device as a microwave oven rather than indiscriminate noise. You can also determine exactly which part of the band is affected by the device and because you can locate it, you can understand which access points are most severely affected. You can then use this information to direct RRM in selecting a channel plan that avoids this source of interference for the access points within its range. Because this interference is not active for a large portion of the day, existing RF management applications might attempt to again change the channels of the affected access points. Persistent device avoidance is unique, however, in that it remains in effect as long as the source of interference is periodically detected to refresh the persistent status. The Cisco CleanAir system knows that the microwave oven exists and includes it in all future planning. If you move either the microwave oven or the surrounding access points, the algorithm updates RRM automatically.

**Note**

Spectrum event-driven RRM can be triggered only by Cisco CleanAir-enabled access points in local mode.

Spontaneous interference is interference that appears suddenly on a network, perhaps jamming a channel or a range of channels completely. The Cisco CleanAir spectrum event-driven RRM feature allows you to set a threshold for air quality (AQ) that, if exceeded, triggers an immediate channel change for the affected access point. Most RF management systems can avoid interference, but this information takes time to propagate through the system. Cisco CleanAir relies on AQ measurements to continuously evaluate the spectrum and can trigger a move within 30 seconds. For example, if an access point detects interference from a video camera, it can recover by changing channels within 30 seconds of the camera becoming active. Cisco CleanAir also identifies and locates the source of interference so that more permanent mitigation of the device can be performed at a later time.

In the case of Bluetooth devices, Cisco CleanAir-enabled access points can detect and report interferences only if the devices are actively transmitting. Bluetooth devices have extensive power save modes. For example, interference can be detected when data or voice is being streamed between the connected devices.

## Persistent Devices

Some interference devices such as outdoor bridges and microwave ovens only transmit when needed. These devices can cause significant interference to the local WLAN due to short duration and periodic operation remain largely undetected by normal RF management metrics. With CleanAir the RRM DCA algorithm can detect, measure, register, and remember the impact and adjust the DCA algorithm. This minimizes the use of channels affected by the persistent devices in the channel plan local to the interference source. Cisco CleanAir detects and stores the persistent device information in the Cisco WLC and this information is used to mitigate interfering channels.

## Persistent Devices Detection

CleanAir-capable Monitor Mode access point collects information about persistent devices on all configured channels and stores the information in the Cisco WLC. Local/Bridge mode AP detects interference devices on the serving channels only.

## Persistent Devices Propagation

Persistent device information that is detected by local or monitor mode access points is propagated to the neighboring access points connected to the same Cisco WLC to provide better chance of handling and avoiding persistent devices. Persistent device detected by the CleanAir-enabled access point is propagated to neighboring non-CleanAir access points, thus enhancing channel selection quality.

## Detecting Interferers by an Access Point

When a CleanAir-enabled access point detects interference devices, detections of the same device from multiple sensors are merged together to create clusters. Each cluster is given a unique ID. Some devices conserve power by limiting the transmit time until actually needed, which results in the spectrum sensor to temporarily stop detecting the device. This device is then correctly marked as down. A down device is correctly removed from the spectrum database. In cases when all the interferer detections for a specific devices are reported, the cluster ID is kept alive for an extended period of time to prevent possible device detection bouncing. If the same device is detected again, it is merged with the original cluster ID and the device detection history is preserved.

For example, some Bluetooth headsets operate on battery power. These devices employ methods to reduce power consumption, such as turning off the transmitter when not actually needed. Such devices can appear to come and go from the classification. To manage these devices, CleanAir keeps the cluster IDs longer and they are remerged into a single record upon detection. This process smoothens the user records and accurately represents the device history.

For more detailed discussion CleanAir and configuration of the same, refer to:

[http://www.cisco.com/c/en/us/td/docs/wireless/controller/7-6/configuration-guide/b\\_cg76/b\\_cg76\\_chapter\\_01000011.html](http://www.cisco.com/c/en/us/td/docs/wireless/controller/7-6/configuration-guide/b_cg76/b_cg76_chapter_01000011.html).

## Post-Deployment RF Tuning

In addition to having both RRM and CleanAir turned on, it is recommended that regular post-deployment RF Site Survey is performed. A post-deployment RF Site Survey ensures that spectrum is being effectively managed by the CMX solution and there are no surprises. Just like the Pre-Deployment RF Site Survey, a Post-Deployment Site Survey should look at the following:

### Location Assessment

- Assess building type and materials used—Has the building under gone renovation with newer material used or newer obstacles erected?
- Areas where full coverage and full performance is needed—Have coverage zones changed? Zones that were thinly populated with APs before because of no coverage requirement may now be required to have full coverage.
- Areas where RF free zones exist—Are there new RF free zones that did not exist before?
- Obtain location maps for RF surveys later—Have any maps been updated with newer construction or layouts.
- High Density—Has this location become a high density location that was not anticipated? In this case a reevaluation of location should be performed.

## Business Needs of WLAN

- **Critical Services**—Are the services still the same? Is the network able to guarantee the same performance as before? Features like AVC/QOS should be explored.
- **Long term Network View**—Planning for future needs is always better than planning for present. With explosion of smartphones and mobile devices in markets, more and more devices are constantly on the Wi-Fi network than ever before. Internet of Everything is enabling machine-to-machine and machine-to-network communication via Wi-Fi. It is not just enough to plan for smartphones and laptops alone, but also for potential IoT devices that may use Wi-Fi.
- **Guarding against potential RF explosion**—Have number of devices gone up way more than anticipated? In this case better (and more) APs may be required leading to a high density design later.
- **802.11n/ac readiness and expectations**—Is the customer looking at upgrading their wireless infrastructure? A new network should be re planned with all new objectives in mind, which may lead to a different RF plan than before.

## Constraints on Deployment

- **DFS and radar avoidance requirement**—Are there any new radar and DFS requirement that did not exist before.
- **Aesthetic design requirements**—Aesthetics requirements may change over time. Administrators should always keep a watch on whether aesthetic requirements over time have changed AP location and antenna direction.
- **Are there newer channels available**—The FCC makes newer spectrum available time to time for WLAN to use. Can these enabled if the software supports it?

## Existing 802.11 Surveys

It is a good idea to capture existing 802.11 state of the location after CMX is deployed and at regular intervals. Because Wi-Fi and other technologies are hugely prevalent today, it is a good investment of time to have a pre-study of existing RF done. Use tools like Metageek site survey to record existing RF at the location. In a big location it is advisable to move through the entire location and take a RF base reading.

- **Plan for persistent non-movable interferers**—Are there newer persistent interferers in the location that were not there before (microwaves, video cameras, etc.)
- **Capture current state of Wi-Fi network**—How many SSIDs exist already? What is their signal strength? Are there potential Wi-Fi networks around the location that are strong enough to interfere with a CMX deployment? What can be done to mitigate such potential sources of interference? Can power be increased on the APs?
- **Antenna Evaluation**—Evaluate use external antennas instead of internal antennas to ensure good coverage based on above parameters. Will use of external antenna be an aesthetic constraint? Can directional antennas be used in corner of the building to direct more coverage into building rather than outside of it? In a large environment this may be useful.—







## Best Practices Checklist

---

**September 4, 2014**

This chapter discusses the best practices check list while deploying a CMX solution.

1. Familiarize yourself with RF basics.
2. Familiarize yourself with 802.11 fundamentals and location fundamentals.
3. Check for regulatory domain restrictions in the area.
4. Conduct pre-site survey:
  - Assess building type and materials used
  - Anticipate difficult zones:
    - Areas where full coverage and full performance is needed.
    - Areas where location is important and needed.
    - Areas where RF free zones exist.
  - Business needs of the WLAN
  - How will the network look in two years.
  - Plan for future client explosion
  - Scale of the planning
  - Expected audience of the network
  - Constraints on deployment
  - DFS and radar avoidance requirement
  - Aesthetic design requirements
  - Predictive surveys for simple budgeting
  - Plan for cabling, power drops, and power requirements
  - Obtain location maps for RF surveys later
  - Thorough versus sample area survey
  - 802.11n/ac readiness and expectations
  - Outdoor readiness
  - Location where access points can be deployed
  - Capture existing 802.11

- Plan for persistent non-movable interferers
  - Antenna evaluation
  - High density consideration
5. Conduct predictive site survey.
- Use Prime Infrastructure RF Planner tool to design a RF plan
  - Use Ekahau Site Survey planner to design a RF plan in absence of Cisco Prime Infrastructure
  - Location tracking = AP per 2500 sqft
  - Enter signal strength to be -67 dBm
  - SNR should to be  $\geq 20$  dBm
  - Data rate to be 12 mbps
  - Ping round trip should be  $\leq 500$ ms
  - Packet loss should be  $\leq 10\%$
  - Remember, minimum number of APs for location tracking is three.
6. AP placement guidelines
- APs should be  $< 70$  feet apart.
  - Use directional antennas where necessary.
  - Plan for coverage with location calculation in mind.
  - Use staggered layout for placing APs, not a straight line.
  - Plan to avoid interferers and installation of APs over interfering sources.
  - Ensure perimeter coverage.
  - Use a single type of AP across the installation.
7. AP Capacity planning
- Plan with client density and client type in mind.
  - Plan for minimum application performance needed.
  - Plan for future growth of client types.
  - Consider high density design if a location demands it.
8. Antennas
- Select the ideal antenna for installation depending on building and aesthetic requirements
  - Be aware of power requirements with antennas.
9. Post-RF deployment
- Turn on RRM and let the network settle before go live.
  - Turn on Clean Air on capable access points.
  - Conduct post-site survey regularly to keep a watch on RF.



## **PART 5**

### **CMX Configuring the Infrastructure**





## Summary of CMX Configuring the Infrastructure

---

**September 4, 2014**

This part of the CVD discusses the configuration of various components necessary for the implementation of a successful CMX solution. These components include Cisco wireless LAN controllers (WLCs), Cisco Prime Infrastructure (PI), and the Cisco Mobility Services Engine (MSE). In addition, the configuration of CMX services, specifically CMX Analytics and CMX Visitor Connect, are discussed.

This part of the CVD includes the following chapters:

- [Chapter 23, “Configuring Cisco Wireless LAN Controllers”](#)—This chapter highlights the configuration of the B2C guest WLAN necessary for providing guest wireless connectivity as discussed in [B2C Guest Access for CMX Visitor Connect](#) in [Chapter 4, “CMX Deployment Models.”](#) Additionally, it provides information about enabling presence on the WLC.
- [Chapter 24, “Configuring Cisco Prime Infrastructure”](#)—This chapter highlights the configuration of Cisco Prime Infrastructure.
- [Chapter 25, “Configuring the Mobility Services Engine for CMX”](#)—This chapter highlights the configuration of the MSE.
- [Chapter 26, “Configuring CMX Analytics”](#)—This chapter highlights the configuration options available for the three main functional areas of CMX Analytics—Dashboard, Analytics, and Reports. With the CMX Analytics Dashboard, adding, modifying, or deleting pages and widgets are discussed. Each of the different types of analysis and their configuration options are discussed within CMX Analytics. Finally, each of the different types of reports and their configuration options are discussed with CMX Reports.
- [Chapter 27, “Configuring CMX Visitor Connect”](#)—This chapter highlights the configuration required to enable CMX Visitor Connect on the MSE.





# Configuring Cisco Wireless LAN Controllers

September 4, 2014

The CMX design guide uses Cisco 5508 Series and Flex 7510 Series Wireless LAN Controllers (WLCs) with version 8.0 code.

To get started, follow the WLC 5508 8.0 configuration design guidelines at:

[http://www.cisco.com/c/en/us/td/docs/wireless/controller/8-0/configuration-guide/b\\_cg80/b\\_cg80\\_chapter\\_010.html](http://www.cisco.com/c/en/us/td/docs/wireless/controller/8-0/configuration-guide/b_cg80/b_cg80_chapter_010.html).



**Note**

The NTP server should be the same one that you will use for the FastLocate feature with the Wireless Security Module (WSM).

To configure basic guest access with a foreign and anchor controller, follow the design guidelines at: <http://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Apr2014/CVD-CampusWirelessLANDesignGuide-APR14.pdf>.

Once the Wireless LAN Controller has been setup, follow the instructions below to configure a wireless network for Visitor Connect.

## WLC Visitor Connect Configuration

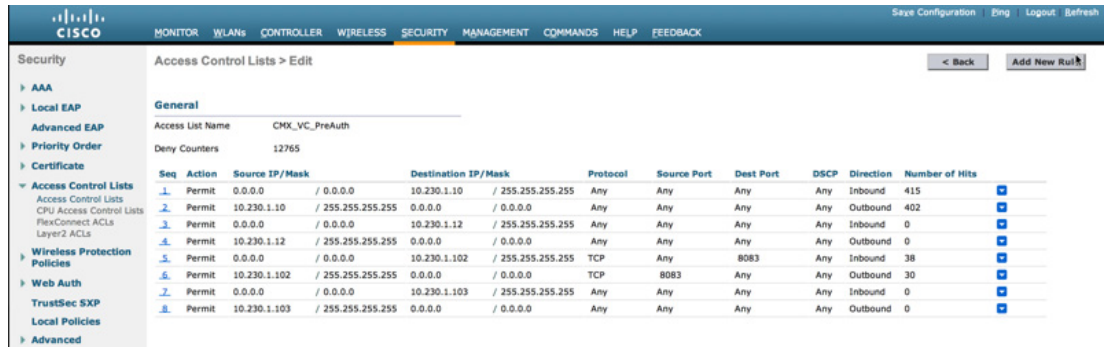
The Visitor Connect configuration on the WLC 5508 consists of two parts:

1. Configuring a pre-authentication ACL on the controller. The pre-authentication ACL causes the wireless LAN controller to redirect all HTTP traffic to the MSE, except traffic which is permitted within the ACL. DNS and DHCP traffic are allowed through in the ACL, as well as traffic destined to the TCP port of the MSE which runs the Visitor Connect service.
2. Configuring Web Passthrough using an external server (pointing to the MSE running CMX Visitor Connect) on the B2C Guest WLAN for Layer 3 security.

## Configuring the ACL for CMX Visitor Connect

Configure the pre-authentication ACL on the Cisco 5508 Series WLC, as shown in [Figure 23-1](#).

Figure 23-1 Configure WLC ACL for CMX Visitor Connect



Seq	Action	Source IP/Mask	Destination IP/Mask	Protocol	Source Port	Dest Port	DSCP	Direction	Number of Hits
1	Permit	0.0.0.0 / 0.0.0.0	10.230.1.10 / 255.255.255.255	Any	Any	Any	Any	Inbound	415
2	Permit	10.230.1.10 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Outbound	402
3	Permit	0.0.0.0 / 0.0.0.0	10.230.1.12 / 255.255.255.255	Any	Any	Any	Any	Inbound	0
4	Permit	10.230.1.12 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Outbound	0
5	Permit	0.0.0.0 / 0.0.0.0	10.230.1.102 / 255.255.255.255	TCP	Any	8083	Any	Inbound	38
6	Permit	10.230.1.102 / 255.255.255.255	0.0.0.0 / 0.0.0.0	TCP	8083	Any	Any	Outbound	30
7	Permit	0.0.0.0 / 0.0.0.0	10.230.1.103 / 255.255.255.255	Any	Any	Any	Any	Inbound	0
8	Permit	10.230.1.103 / 255.255.255.255	0.0.0.0 / 0.0.0.0	Any	Any	Any	Any	Outbound	0

CMX Visitor Connect with Splash Pages & Social Connectors uses port 8083 on the MSE. The example pre-authentication ACL permits (does not redirect) traffic from any host destined for TCP port 8083 of the IP address of the MSE server. Implicitly any traffic not specifically allowed (permitted) is redirected to the MSE running CMX Visitor Connect. This prevents web session traffic already destined for the MSE running CMX Visitor Connect from being re-directed by the wireless LAN controller. In addition DNS and DHCP traffic are typically not redirected via additional access control entries (ACEs) within the pre-authentication ACL.

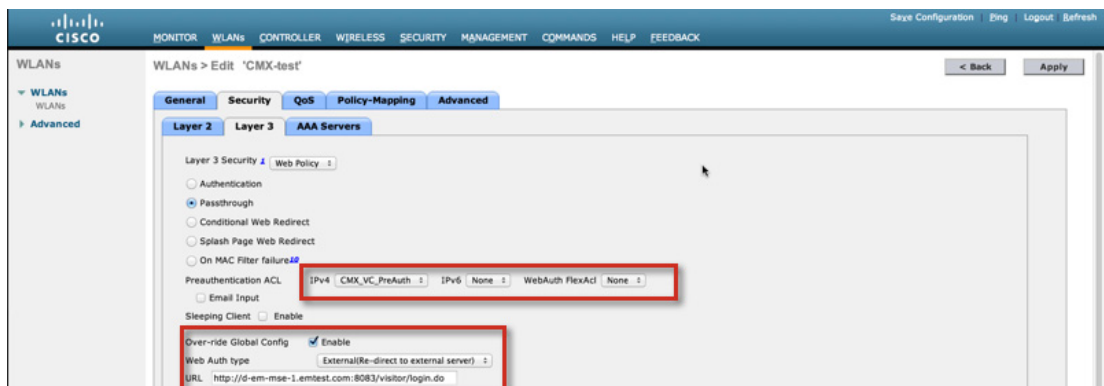
**Note**

In the design presented in this guide, guest traffic is terminated on a dedicated guest wireless LAN controller sitting on a DMZ segment of an ASA firewall. The ASA firewall security policy should also be configured to only allow guests who are using CMX Visitor Connect to access TCP port 8083 of the MSE which runs CMX Visitor connect.

## Configuring the WLAN for Visitor Connect

Figure 23-2 shows an example of the configuration of Web Passthrough using re-direction to an external server for the guest WLAN for CMX Visitor Connect. Additionally, Figure 23-2 shows the application of the pre-authentication ACL to the guest WLAN.

Figure 23-2 Configuring the B2C Guest WLAN for CMX Visitor Connect with Splash Pages



To configure Web Passthrough Using Re-direction for CMX Visitor Connect:



- 
- Step 1** On the Security Layer 3 tab, from the drop-down menu next to Layer 3 Security, select **Web Policy**.
- Step 2** Select the **Passthrough** checkbox as the type of web policy.
- Step 3** Select the name of the IPv4 Pre-authentication ACL configured in the previous section from the drop-down menu next to Pre-authentication ACL.
- Step 4** In the drop-down menu next to Web Auth type, select **External** (Re-direct to external server).
- Step 5** Configure the External Redirect URL to point to the URL:  
http://<IP\_Address\_or\_Name\_of\_MSE\_Server\_running\_Visitor\_Connect>:8083/visitor/login.do  
This redirects web traffic to the CMX Visitor Connect service running on TCP port 8083 of the MSE server.
- Step 6** Click the **Apply** button at the top right corner of the page to apply and save the changes.
- 

## Configuring FastLocate

FastLocate requires WSM modules to be installed on all access points deployed within the site. This section discusses enabling FastLocate directly on a Cisco 5508 Series WLC. Note that you can also use Cisco Prime Infrastructure to enable FastLocate by using templates to enable it on multiple Wireless LAN Controllers.

To configure FastLocate:

- 
- Step 1** Ensure that the switches to which the APs with WSM modules will be connected support POE+ or Enhanced POE.
- Step 2** Ensure that a valid NTP source has been configured on the WLC.
- Step 3** Insert WSM modules into the APs and attach them to POE+ or Enhanced POE switch ports.



---

**Note** The reference Installation for the WSM is available at:  
<http://www.cisco.com/c/en/us/support/docs/wireless/aironet-3600-series/115612-Aironet-Access-Point-Module-for-WSSI-Guide-00.html>

---

- Step 4** Establish an HTTPS session to the WLC which controls the APs. Navigate to **Wireless > Access Points > Global Configuration**. Within the Global Configuration page, scroll down to the Packet RSSI Config. Parameters section. An example is shown in [Figure 23-3](#).

Figure 23-3 Configuring FastLocate

The screenshot shows the Cisco WLC configuration interface. The 'WIRELESS' tab is active. In the left sidebar, 'Global Configuration' is selected. The main content area shows various configuration sections. A red box highlights the 'Packet RSSI Location Config Parameters' section, which includes the following settings:

- Enable Packet RSSI Location:
- Packet Detection RSSI Minimum (dBm): -100
- Scan Count Threshold for Idle Client Detection (dBm): 10
- NTP Server: 172.19.36.1

Other visible sections include:

- 802.1x Supplicant Credentials: 802.1x Authentication (unchecked)
- AP Fallover Priority: Global AP Fallover Priority (Disable)
- AP Image Pre-download: Download Primary, Download Backup, Interchange Image, Abort Predownload
- OEAP Config Parameters: Disable Local Access (unchecked), Enable Split Tunnel (unchecked)
- Flexconnect Ethernet Falback: Radio Interface Shutdown (unchecked), Flexconnect Arp-Cache (unchecked)
- Global Telnet SSH: Telnet (unchecked), SSH (unchecked)
- Global IPv6 UDP Lite: UDP Lite (checked)

Footnote 3 at the bottom states: 'Packet RSSI Location config parameters are applicable only to 3600/3700 APs with WSSI module.'

- Step 5** Click the checkbox next to **Enable Packet RSSI Location**. Leave other details at their default. Configure the IP address of the NTP server that is accessible by the WLC and APs.



**Note** The Scan Count Threshold for Idle Client Detection (dBm) field represents the number of off-channel scan cycles the AP waits before sending a Block Acknowledgement Request (i.e., BAR) to idle clients. The default value of 10 corresponds to approximately 40-60 seconds, depending on the number of channels in the off-channel scan cycle and whether Cisco CleanAir® has been enabled.

- Step 6** Click **Apply** to save the changes and enable FastLocate.



**Note** FastLocate should be enabled on all WLC. It is possible to do this via Templates using Prime Infrastructure 2.1, however it is not described in this guide.



## Configuring Cisco Prime Infrastructure

---

September 4, 2014

The CMX design guide uses Cisco Prime Infrastructure 2.1 as the network management solution to configure, manage, and synchronize Cisco WLC 5508s with the Cisco Mobility Services Engine.

### Installing Cisco Prime Infrastructure

To get started, follow the design guidelines to install Prime Infrastructure 2.1. For the purposes of this design guide, Prime Infrastructure has been installed on a Cisco UCS system with VMware ESXi.

Before installing Cisco Prime Infrastructure, ensure that the system requirements at this URL have been met:

[http://www.cisco.com/c/en/us/td/docs/net\\_mgmt/prime/infrastructure/2-1/quickstart/guide/cpi\\_qsg.html#99048](http://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/infrastructure/2-1/quickstart/guide/cpi_qsg.html#99048).

To Install Cisco Prime Infrastructure on a Virtual Machine, follow the instructions at:

[http://www.cisco.com/c/en/us/td/docs/net\\_mgmt/prime/infrastructure/2-1/quickstart/guide/cpi\\_qsg.html#pgfId-63672](http://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/infrastructure/2-1/quickstart/guide/cpi_qsg.html#pgfId-63672).

### Installing the Cisco Mobility Services Engine

The CMX design guide uses the Cisco Mobility Services Engine (MSE) version 8.0 as the backend for the Cisco CMX solution. To get started, follow the design guidelines to install Cisco MSE 8.0 described in the referenced documents below. For the purposes of this validated design guide, the MSE has been installed on a Cisco UCS with VMware ESXi 5.1.

[http://www.cisco.com/c/en/us/td/docs/wireless/mse/7-6/Virtual\\_appliance/Installation\\_Config\\_Guide/Cisco\\_MSE\\_VA\\_Config\\_Guide.html](http://www.cisco.com/c/en/us/td/docs/wireless/mse/7-6/Virtual_appliance/Installation_Config_Guide/Cisco_MSE_VA_Config_Guide.html)

Once the basic Wireless LAN infrastructure—the Wireless LAN Controllers, Cisco Prime Infrastructure, and the Mobility Services Engine—are up and running, follow the guidelines below to configure the Cisco Connected Mobile Experiences (CMX) solution.

### Adding Wireless LAN Controllers to Cisco Prime Infrastructure

Prime Infrastructure provides two different graphical user interfaces:

- Lifecycle view is organized according to home, design, deploy, operate, report, and administer menus.
- Classic view closely corresponds to the graphical user interface in Cisco Prime Network Control System 1.1 or Cisco Wireless Control System (WCS).

You can switch back and forth between the views by clicking the downward arrow next to your login name in Prime Infrastructure.

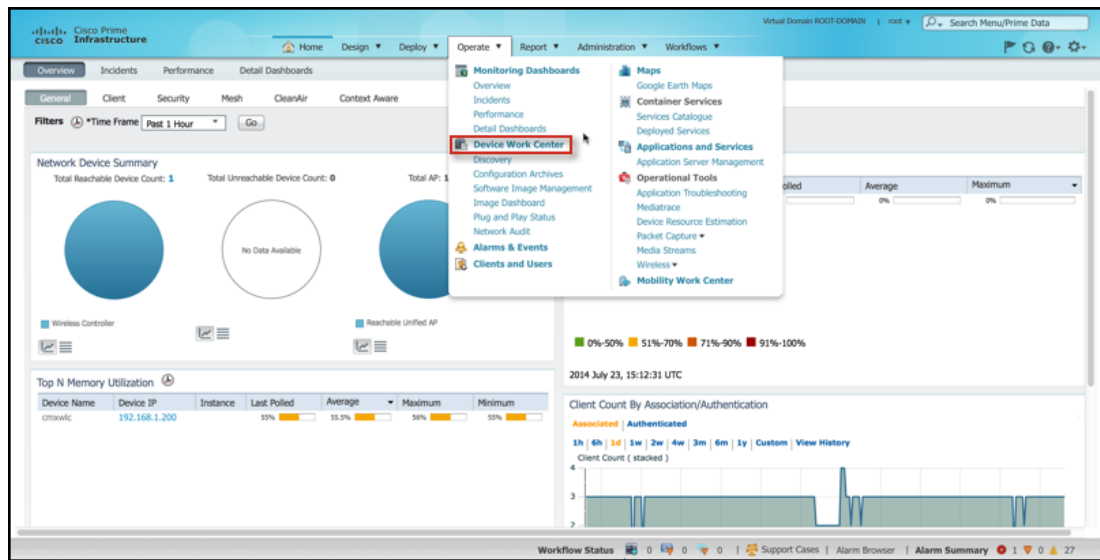


**Note** This design guide uses the Lifecycle view.

To add wireless LAN controllers to Cisco Prime Infrastructure:

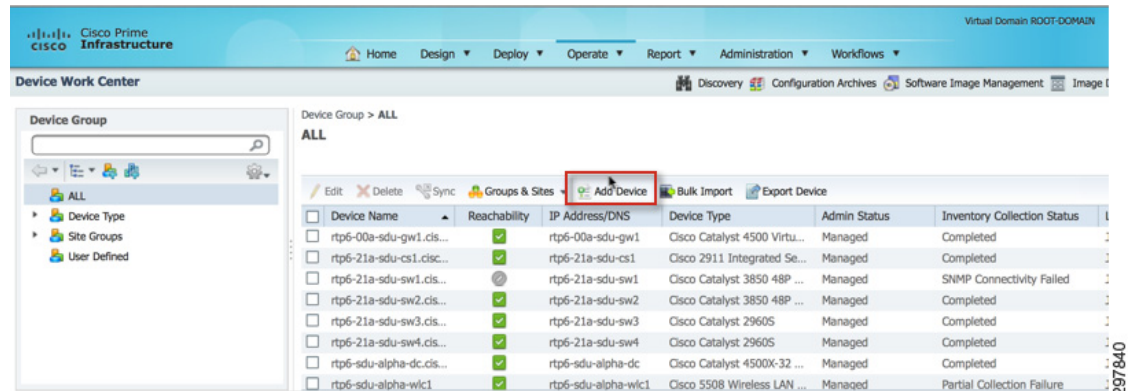
**Step 1** Go to the **Operate** tab and select **Device Work Center**, as shown in [Figure 24-1](#).

**Figure 24-1** Accessing the Device Work Center



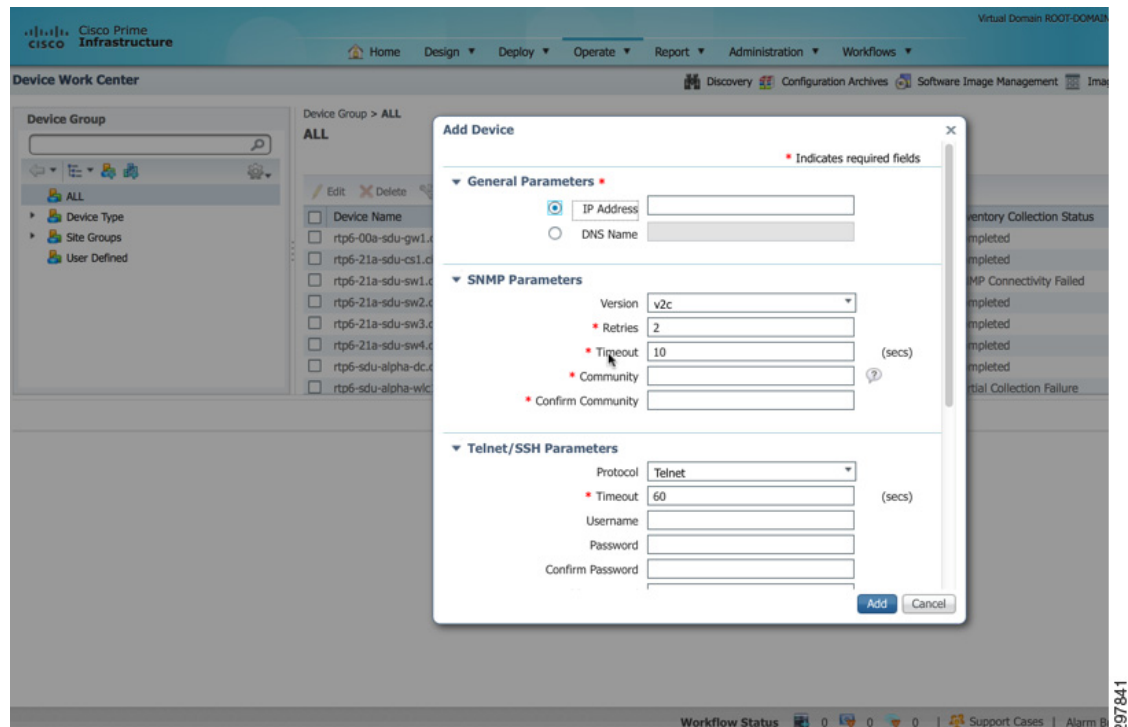
**Step 2** From the Device Work Center page, select **Add Device** from the horizontal menu bar which runs across the center panel of the screen, as shown in [Figure 24-2](#).

Figure 24-2 Add a Device



A popup window similar to the one shown in Figure 24-3 is displayed.

Figure 24-3 Popup Window for Adding a Device



Four sets of parameters regarding the device are requested within the popup window for adding the device:

- **General Parameters**—Add the IP address or DNS name (assuming the device has been added to the DNS server) of the device so Cisco Prime Infrastructure knows how to reach the device.
- **SNMP Parameters**—If the device supports SNMP, choose the version: SNMP v1, 2c, or 3. When selecting SNMP v1 or 2c, configure the SNMP community string. Cisco Prime Infrastructure can be used for monitoring purposes only by adding the read only SNMP community string. If Cisco Prime Infrastructure is used for configuration purposes, use the read/write SNMP community string. Note that SNMP v2c sends the community string in clear text across the network. Hence, an alternative

for configuration purposes is to select SNMP v3 if supported by the device and Cisco Prime Infrastructure since it provides a higher level of security. SNMP v3 configuration requires a username, authentication password, and privacy password. In addition, the authentication type—none, HMAC-MD5, or HMAC-SHA—and the privacy type—none, CBC-DES, or CBF-AES-128—are required. Cisco 5508 and Flex 7500 Series controllers support SNMP v1, v2c, and v3.

- **Telnet/SSH Parameters**—If the device supports Telnet or SSH v2 access, choose one. Note that Telnet sends passwords in clear text across the network. SSH v2 provides a higher level of security since it encrypts the transport. Configure the username and password. Some devices, such as Cisco routers and switches, also require an enable password for configuration purposes. Configure an enable password if the device supports one. Cisco 5508 and Flex 7500 Series controllers support both Telnet and SSH v2 and do not require an enable password.
- **HTTP Parameters**—If the device supports HTTP or HTTPS access, choose one. HTTPS provides a higher level of security since it encrypts the transport. Configure the username and password. Some devices may also support read only access via HTTP or HTTPS. For these devices, a monitor username and password can be supplied if Cisco Prime Infrastructure is to be used for monitoring purposes only. Cisco 5508 and Flex 7500 Series controllers support both read only and read/write HTTP and HTTPS access.

For each of the parameters configured above, it is assumed that the WLC has been previously configured the same way.

Once the WLC has been added, Prime Infrastructure attempts to connect to the device to discover what type of device it is and take inventory of the device and its configuration. The WLC should appear within the list of devices with a status of managed and reachable, as shown in [Figure 24-4](#).

**Figure 24-4** Example of a Wireless LAN Controller Added to Cisco Prime Infrastructure

The screenshot displays the Cisco Prime Infrastructure Device Work Center interface. The top navigation bar includes Home, Design, Deploy, Operate, Report, Administration, and Workflows. The main content area shows a list of devices under the 'ALL' group. A table lists various devices, with the 'rtp6-sdu-alpha-wlc1' device highlighted in red. Below the table, the 'Device Details' section for 'rtp6-sdu-alpha-wlc1' is shown, including a summary of the device type (Cisco 5508 Series Wireless Controller), IP address (10.122.78.26), and name (rtp6-sdu-alpha-wlc1).

Device Name	Reachability	IP Address/DNS	Device Type	Admin Status	Inventory Collection Status
rtp6-00a-sdu-gw1.cis...	✓	rtp6-00a-sdu-gw1	Cisco Catalyst 4500 Virtu...	Managed	Completed
rtp6-21a-sdu-cs1.cis...	✓	rtp6-21a-sdu-cs1	Cisco 2911 Integrated Se...	Managed	Completed
rtp6-21a-sdu-sw1.cis...	⊗	rtp6-21a-sdu-sw1	Cisco Catalyst 3850 48P ...	Managed	SNMP Connectivity Failed
rtp6-21a-sdu-sw2.cis...	✓	rtp6-21a-sdu-sw2	Cisco Catalyst 3850 48P ...	Managed	Completed
rtp6-21a-sdu-sw3.cis...	✓	rtp6-21a-sdu-sw3	Cisco Catalyst 2960S	Managed	Completed
rtp6-21a-sdu-sw4.cis...	✓	rtp6-21a-sdu-sw4	Cisco Catalyst 2960S	Managed	Completed
rtp6-sdu-alpha-dc.cis...	✓	rtp6-sdu-alpha-dc	Cisco Catalyst 4500X-32 ...	Managed	Completed
rtp6-sdu-alpha-wlc1	✓	rtp6-sdu-alpha-wlc1	Cisco 5508 Wireless LAN ...	Managed	Partial Collection Failure



# Configuring Maps within Cisco Prime Infrastructure

To configure and add maps for your deployment, follow the guidelines linked below. Once the buildings and maps have been added, continue with CMX specific requirements to set up maps.

Before getting started it is a good idea to review [Chapter 10, “Radio Frequency Fundamentals.”](#)

1. Ensure that you have maps for all the buildings and floors available to you.
2. Ensure that maps are to scale in Prime Infrastructure.
3. To place APs on the floor, follow the RF plan that you have generated with a combination of the RF planner tool in Prime Infrastructure (or the Ekahau Site Survey tool) and physical RF Site Survey.
4. Ensure you have GPS markers for all the zones.
5. Coverage areas should be configured via Prime Infrastructure. These are mapped to zones within CMX solution.
6. Ensure that Inclusion and Exclusion areas are configured via the Prime Infrastructure Tool.

## Adding Floor Areas to a Campus Building or a Standalone Building

This section describes how to add floor plans to either a campus building or a standalone building in the Prime Infrastructure database. After you add a building to a campus map, you can add individual floor plan and basement maps to the building.

**Note**

Use the zoom controls at the top of the campus image to increase or decrease the size of the map view and to hide or show the map grid (which shows the map size in feet or meters).

To add a floor area to a campus building, create a building under a campus, and then add floor areas to the building.

**Step 1**

Save your floor plan maps in .PNG, .JPG, .JPEG, or .GIF format.

**Note**

For CMX, it is recommended that the size of image file is a maximum of 500 Kbytes. Loading large images into the 3D version of CMX causes certain browsers to show black images. The mse.properties file can also be configured to automatically compress the image.

**Note**

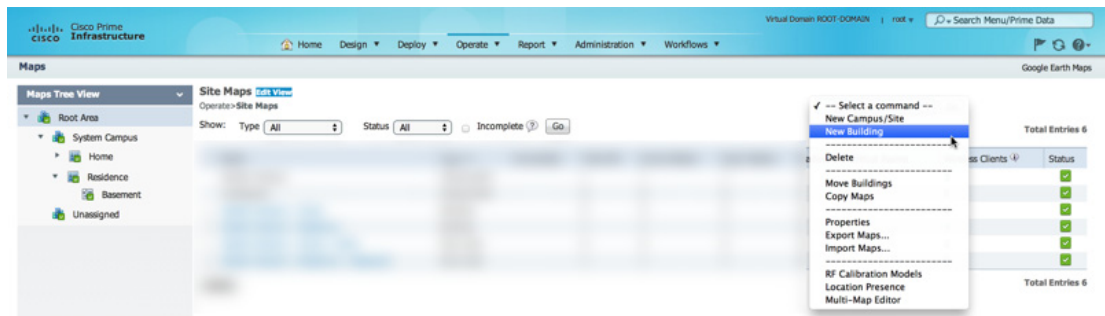
If there are problems converting the auto-cad file, an error message is displayed. The Prime Infrastructure uses a native image conversion library to convert auto-cad files into raster formats like .png. If the native library cannot be loaded, the Prime Infrastructure shows an “unable to convert the auto-cad file” message. If you receive this error, make sure all the required dependencies are met for the native library. To find any dependency problems, use ldd on Linux platforms. The following DLLs must be present under the /webnms/rfdlls Prime Infrastructure installation directory: LIBGFL254.DLL, MFC71.DLL, MSVCR71.DLL, and MSVCP71.DLL. If dependency problems occurs, you have to install the required libraries and restart Prime Infrastructure.

The floor map image is enhanced for zooming and panning. The floor image is not visible completely until this operation is complete. You can zoom in and out to view the complete map image. For example, if you have a high resolution image (near 181 megapixels) whose size is approximately 60 megabytes, it may take two minutes to appear on the map.

**Step 2** Choose **Operate > Maps**.

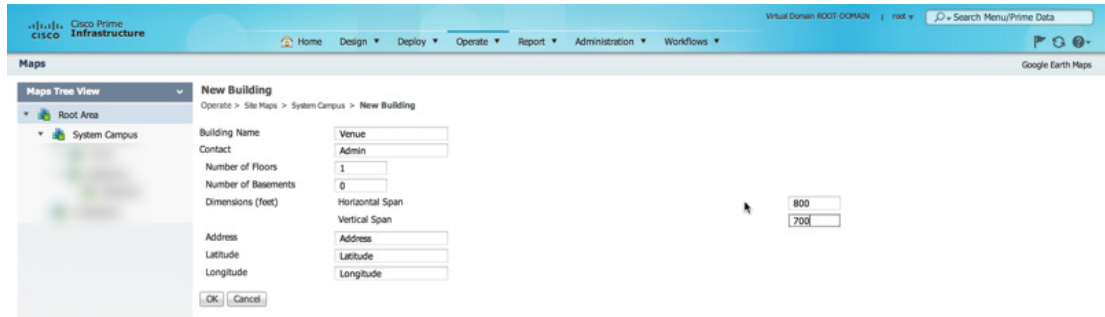
**Step 3** From the Maps Tree View or the **Monitor > Site Maps** list, choose the applicable campus building to open the Building View page. If no Building exists, select **Add Building** from the command drop-down list from the top right and add a building, as shown in [Figure 24-5](#).

**Figure 24-5** Adding a New Building to Prime Infrastructure



Enter building details as shown [Figure 24-6](#) and click **OK**.

**Figure 24-6** Building Details



Continue with Step 4 by clicking the **building name** under Campus.


**Step 4** Hover your mouse cursor over the name within an existing building rectangle to highlight it.

You can also access the building from the Campus View page. In the Campus View page, click the **building name** to open the Building View page.

**Step 5** From the **Select a command** drop-down list, choose **New Floor Area**, as shown in [Figure 24-7](#).



Figure 24-7 Adding a New Floor Area



**Step 6** Click **Go**. The New Floor Area page is displayed, as shown in Figure 24-8.

Figure 24-8 New Floor Area Page



**Step 7** In the New Floor Area page, follow these steps to add floors to a building to organize related floor plan maps:

- a. Enter the **Floor Area Name** and **Contact** information.
- b. Choose the floor or basement number from the **Floor** drop-down list.
- c. Choose the **RF Model** for the floor from the **Floor-Type** drop-down list. Cisco Prime Infrastructure ships with several default RF models that facilitate setup under several common environments:
  - Cubes and Walled Offices
  - Drywall Offices Only
  - Outdoor Open Space
  - Indoor High Ceiling

Choose the one that most closely matches the environment of your floor.

- d. Enter the floor-to-floor height in feet.  
To change the unit of measurement (feet or meters), choose **Monitor > Site Maps** and choose **Properties** from the **Select a command** drop-down list.
- e. Browse to and choose the desired floor or basement image or CAD filename and click **Open**.  
If you are importing a CAD file, use the **Convert CAD File** drop-down list to determine the image file for conversion.



**Tip** It is not recommended to use a .JPEG (.JPG) format for an auto-cad conversion. Unless a JPEG is specifically required, use .PNG or .GIF format for higher quality images.

- f. Click **Next**. At this point, if a CAD file was specified, a default image preview is generated and loaded.




---

**Note** The Prime Infrastructure uses a native image conversion library to convert auto-cad files into raster formats like .PNG. When there are issues loading the native library, Prime Infrastructure shows the following error: “Unable to convert the auto-cad file. Reason: Error while loading the auto-cad image conversion library.” For more information see Prime Infrastructure online help or Prime Infrastructure documentation.

---

The names of the CAD file layers are listed with check boxes to the right side of the image indicating which are enabled.

When you choose the floor or basement image filename, the Prime Infrastructure shows the image in the building-sized grid.

The maps can be of any size because the Prime Infrastructure automatically resizes the maps to fit the workspace.

- g. If you have CAD file layers, you can select or deselect as many as you want and click **Preview** to view an updated image. Click **Next** when you are ready to proceed with the selected layers.

Enter the remaining parameters for the floor area.

- h. Either leave the **Maintain Aspect Ratio** check box selected to preserve the original image aspect ratio or unselect the check box to change the image aspect ratio.
- i. Enter an approximate floor or basement horizontal and vertical span (width and depth on the map) in feet.

The horizontal and vertical spans should be smaller than or the same size as the building horizontal and vertical spans in the Prime Infrastructure database.

- j. If applicable, enter the horizontal position (distance from the corner of the outdoor area rectangle to the left edge of the campus map) and vertical position (distance from the corner of the outdoor area rectangle to the top edge of the campus map) in feet or meters.




---

**Tip** Use **Ctrl-click** to resize the image within the building-sized grid.

---

- k. If desired, select the **Launch Map Editor after floor creation** check box to rescale the floor and draw walls.
- l. Click **OK** to save this floor plan to the database. The floor is added to the Maps Tree View and the Design > Site Maps list.

Use different floor names in each building. If you are adding more than one building to the campus map, do not use a floor name that exists in another building. This overlap causes incorrect mapping information between a floor and a building.

**Step 8** Click any of the floor or basement images to view the floor plan or basement map.

You can zoom in or out to view the map at different sizes and you can add access points.

---

## Adding APs on Maps

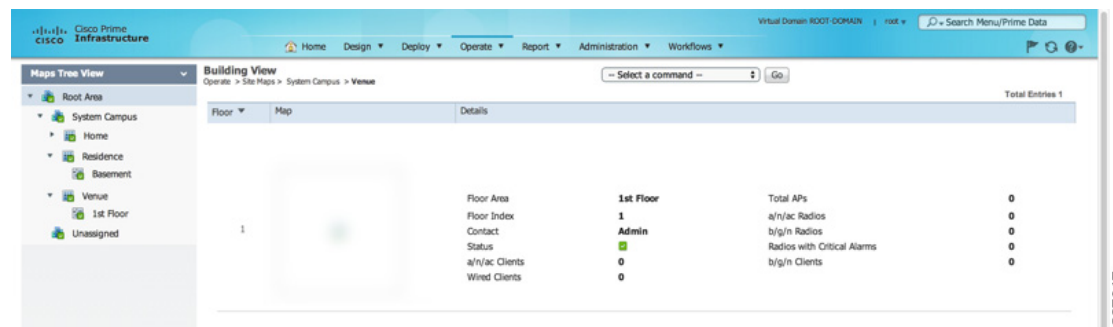
Based on the RF Plan model you have generated by using a combination of physical RF site survey and predictive site survey tools like the Prime Infrastructure RF Planner tool (or Ekahau Site Survey), place the APs on the maps at the desired location. To place APs on the maps, follow the steps in the sections below.

### Adding Access Points to a Floor Area

After you add the .PNG, .JPG, .JPEG, or .GIF format floor plan and outdoor area maps to the Prime Infrastructure database, you can position lightweight access point icons on the maps to show where they are installed in the buildings. To add access points to a floor area and outdoor area, follow these steps:

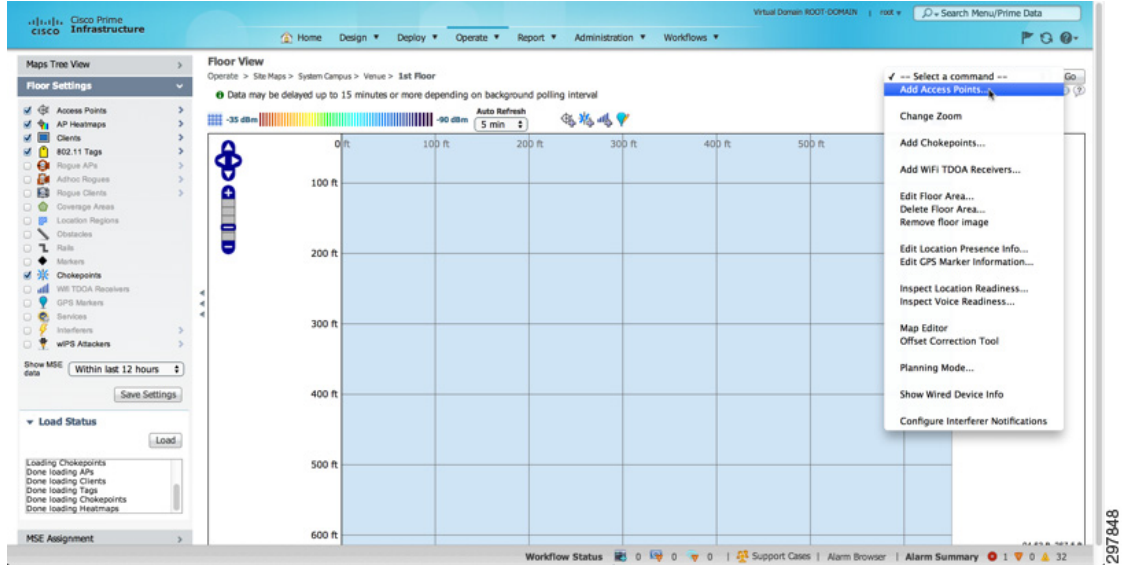
- Step 1** Choose **Operate > Maps** to display the Maps page. The Maps Tree View in the left panel of the page can be expanded to show the various campuses, buildings, and floors defined in Prime Infrastructure. Selecting one of the buildings displays a Building View page, as shown in [Figure 24-9](#).

**Figure 24-9** Building View Page



- Step 2** From the Maps Tree View in the left panel of the page, select the applicable floor within the building to open the Floor View page, as shown in [Figure 24-10](#).

Figure 24-10 Floor View Page



**Step 3** From the Select a command drop-down list on the right side of the page, choose **Add Access Points** and click **Go**.

**Step 4** In the Add Access Points page, select the check boxes of the access points that you want to add to the floor area.

**Note**

If you want to search for access points, enter AP name or MAC address (Ethernet/Radio)/IP in the Search AP [Name/MacAddress (Ethernet/Radio)/IP] text box and then click **Search**. The search is case-insensitive.

**Note**

Only access points that are not yet assigned to any floor or outdoor area appear in the list.

**Note**

Select the check box at the top of the list to select all access points.

**Step 5** When all of the applicable access points are selected, click **OK** at the bottom of the access point list. The Position Access Points page is displayed.

Each access point you have chosen to add to the floor map is represented by a gray circle (differentiated by access point name or MAC address) and is lined up in the upper left part of the floor map.

**Step 6** Click and drag each access point to the appropriate location. Access points turn blue when selected.

**Note**

When you drag an access point on the map, its horizontal and vertical position appears in the Horizontal and Vertical text boxes.

**Note**

The small black arrow at the side of each access point represents Side A of each access point and each access point arrow must correspond with the direction in which the access points were installed. Side A is clearly noted on each 1000 series access point and has no relevance to the 802.11a/n radio. To adjust the directional arrow, choose the appropriate orientation from the Antenna Angle drop-down list.

When selected, the access point details are displayed on the left side of the page and include:

- AP Model—Indicates the model type of the selected access point.
- Protocol—Choose the protocol for this access point from the drop-down list.
- Antenna—Choose the appropriate antenna type for this access point from the drop-down list.
- Antenna/AP Image—The antenna image reflects the antenna selected from the Antenna drop-down list. Click the arrow at the top right of the antenna image to expand the image size.
- Antenna Orientation—Depending on the antenna type, enter the Azimuth and the Elevation orientations in degrees.

**Note**

The Azimuth option does not appear for Omnidirectional antennas because their pattern is nondirectional in azimuth.

**Note**

For internal antennas, the same elevation angle applies to both radios.

The antenna angle is relative to the map X axis. Because the origin of the X (horizontal) and Y (vertical) axes is in the upper left corner of the map, 0 degrees points side A of the access point to the right, 90 degrees points side A down, 180 degrees points side A to the left, and so on.

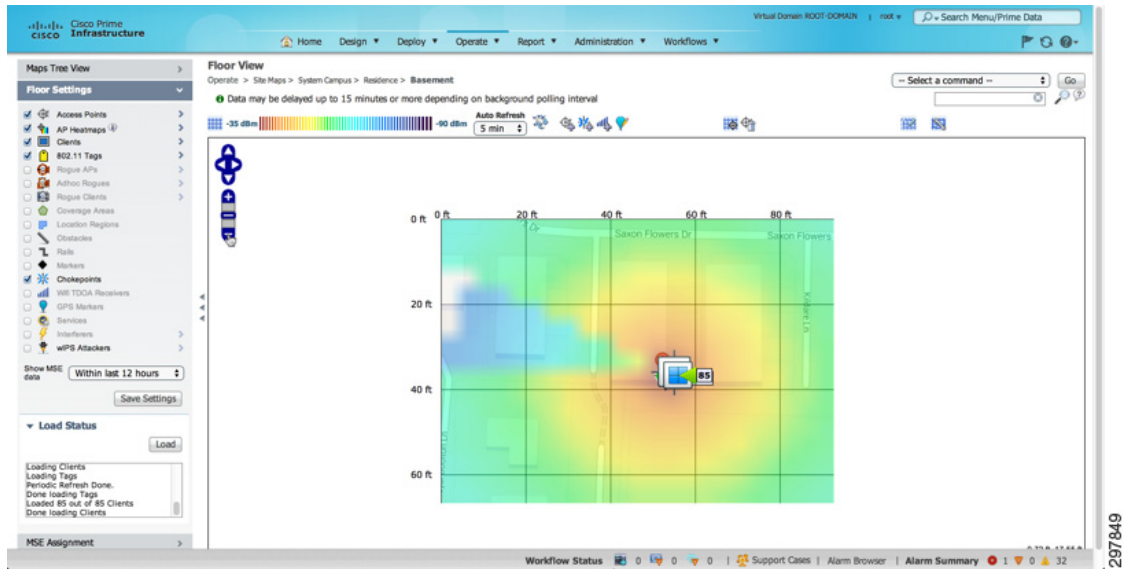
The antenna elevation is used to move the antenna vertically, up or down, to a maximum of 90 degrees.

**Note**

Make sure each access point is in the correct location on the map and has the correct antenna orientation. Accurate access point positioning is critical when you use the maps to find coverage holes and rogue access points.

See the following URL for further information about the antenna elevation and azimuth patterns:  
[http://www.cisco.com/en/US/products/hw/wireless/ps469/tsd\\_products\\_support\\_series\\_home.html](http://www.cisco.com/en/US/products/hw/wireless/ps469/tsd_products_support_series_home.html).

Figure 24-11 Floor View with Map



**Step 7** When you are finished placing and adjusting each access point, click **Save**.



**Note**

Clicking **Save** causes the antenna gain on the access point to correspond to the selected antenna, which might cause the radio to reset.

Prime Infrastructure computes the RF prediction for the coverage area. These RF predictions are popularly known as heat maps because they show the relative intensity of the RF signals on the coverage area map.



**Note**

This display is only an approximation of the actual RF signal intensity because it does not take into account the attenuation of various building materials, such as drywall or metal objects, nor does it display the effects of RF signals bouncing off obstructions.



**Note**

Antenna gain settings have no effect on heatmaps and location calculations. Antenna gain is implicitly associated to the antenna name. Because of this, the following apply:

- If an antenna is used and marked as “Other” in Prime Infrastructure, it is ignored for all heatmap and location calculations.
- If an antenna is used and marked as a Cisco antenna in the Prime Infrastructure, that antenna gain setting (internal value on Prime Infrastructure) is used no matter what gain is set on the controller.

## Defining Coverage Area

To draw a coverage area using the Prime Infrastructure UI, follow these steps:



**Note**

You must add floor plan before drawing a coverage area.

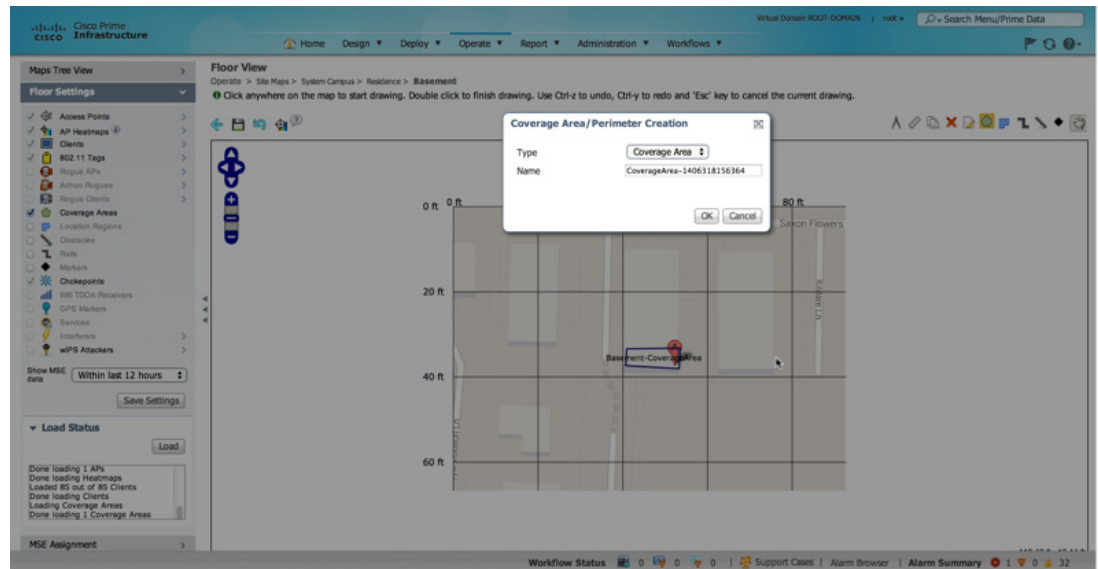
- Step 1** Add the floor plan if it is not already represented in the Prime Infrastructure.
- Step 2** Choose **Operate > Maps**.
- Step 3** Click the **Map Name** that corresponds to the outdoor area, campus, building, or floor you want to edit.
- Step 4** From the **Select a command** drop-down list, choose **Map Editor**, and click **Go**.
- Step 5** In the **Map Editor** page, click the **Draw Coverage Area** icon on the toolbar. The coverage area icon is shown as:



A pop-up is displayed.

- Step 6** Enter the name of the area that you are defining and click **OK**.
- A drawing tool is displayed.
- Step 7** Move the drawing tool to the area you want to outline.
- Click the left mouse button to begin and end drawing a line.
  - When you have completely outlined the area, double-click the left mouse button and the area is highlighted in the page.
- The outlined area must be a closed object to appear highlighted on the map.
- Step 8** Click the **disk icon** on the toolbar to save the newly drawn area.

**Figure 24-12** Configure Coverage Area



**Note**

Coverage Zones cannot be renamed in this version of MSE software release.

## Monitoring Geo-Location

The MSE provides the physical location of wired clients, wired endpoints, switches, controllers, and access points present in a wireless network deployment. Currently MSE provides location information in geo-location format to the external entities through northbound and southbound entities.

To improve the accuracy of the geo-location information provided by MSE, this feature aims to transform the geometric location co-ordinates of a device to geo-location coordinates (latitude and longitude) and provides it to the external entities through northbound and southbound interfaces.




---

**Note** At least three GPS markers are required for geo-location calculation. The maximum number of GPS markers that you can add is 20.

---




---

**Note** For CMX Analytics, the 2D OpenStreetMaps requires all points to be geo-located as latitude/longitude for the results to be displayed in the correct geographical location.

---

## Adding a GPS Marker to a Floor Map

To add a GPS marker to a floor map, follow these steps:

- 
- Step 1** Choose **Operate > Maps** to display the Maps page.
  - Step 2** Choose **Campus Name > Building Name > Floor Name**.
  - Step 3** Choose the **Add/Edit GPS Markers Information** menu option on the top left menu to open the Add/Edit GPS page.

A GPS Marker icon is displayed on the top left corner of the map (X=0 Y=0).

- Step 4** You can drag the GPS Marker icon and place it in the desired location on the map or enter the X and Y position values in the GPS Marker Details table on the left sidebar menu to move the marker to the desired position.




---

**Note** If the markers added are too close, then the accuracy of geo-location information is less.

---

- Step 5** Enter the Latitude and Longitude degrees for the selected GPS Marker icon in the left sidebar menu and click **Save**.

The GPS Marker information is saved to the database.

- Step 6** Click **Apply to other Floors of Building** to copy GPS markers on one floor of a building to all the remaining floors of that building.




---

**Note** The GPS marker information is required by the CMX analytics to show results for the building in the 2D Open Street Maps view. A warning message is displayed if these GPS markers are not set. The latitude or longitude of the GPS markers can often be obtained through mapping software such as Google maps or Open Street Maps.

---



## Editing a GPS Marker

To edit a GPS marker present on the floor, follow these steps:

- 
- Step 1** Choose **Operate > Maps** to display the Maps page.
  - Step 2** Choose the **Campus Name > Building Name > Floor Name**.
  - Step 3** Choose the **Add/Edit GPS Markers Information** menu option on the top left menu to open the Add/Edit GPS page.
  - Step 4** Select an existing GPS marker present on the floor.
  - Step 5** From the left sidebar menu, you can change the Latitude, Longitude, X Position, and Y Position which is associated with the GPS marker.
  - Step 6** Click **Save**.
- The modified GPS marker information is now saved to the database.
- 

## Deleting a GPS Marker Present on a Floor

To delete a GPS marker present on a floor, follow these steps:

- 
- Step 1** Choose **Operate > Maps** to display the Maps page.
  - Step 2** Choose **Campus Name > Building Name > Floor Name**.
  - Step 3** Choose the **Add/Edit GPS Markers Information** menu option to open the Add/Edit GPS page.
  - Step 4** Select an existing GPS Marker which is present on the floor from the left sidebar menu.



---

**Note** You can delete multiple GPS markers present on a floor by selecting the **Multiple GPS Markers** check box.

---

- Step 5** Click **Delete GPS Marker**.
- The selected GPS marker is deleted from the database.
- 

## Inclusion and Exclusion Areas on a Floor

- Inclusion and exclusion areas can be any polygon shape and must have at least three points. Points can sometime be located outside the building. If this is where the devices are, then a coverage area should be created. At other times, the points are actually inside and should be moved to the nearest inside location (the same applies for unlikely areas inside). Defining inclusion and exclusion areas does this and therefore the analytic results are more consistent.
- You can only define one inclusion region on a floor. By default, an inclusion region is defined for each floor when it is added to Prime Infrastructure. The inclusion region is indicated by a solid aqua line and generally outlines the region.
- You can define multiple exclusion regions on a floor.

- Newly defined inclusion and exclusion regions appear on heatmaps only after the mobility services engine recalculates location.

## Defining an Inclusion Region on a Floor

To define an inclusion area, follow these steps:

- 
- Step 1** Choose **Operate > Maps**.
  - Step 2** Click the name of the appropriate floor area.
  - Step 3** From the **Select a command** drop-down list, choose **Map Editor**.
  - Step 4** Click **Go**.
  - Step 5** At the map, click the **aqua box** on the toolbar.




---

**Note** A message box appears reminding you that only one inclusion area can be defined at a time. Defining a new inclusion region automatically removes the previously defined inclusion region. By default, an inclusion region is defined for each floor when it is added to Prime Infrastructure. The inclusion region is indicated by a solid aqua line and generally outlines the region.

---

- Step 6** Click **OK** in the message box that appears. A drawing icon appears to outline the inclusion area.
- Step 7** To begin defining the inclusion area, move the drawing icon to a starting point on the map and click once.
- Step 8** Move the cursor along the boundary of the area you want to include and click to end a border line. Click again to define the next boundary line.
- Step 9** Repeat Step 8 until the area is outlined and then double-click the **drawing icon**. A solid aqua line defines the inclusion area.
- Step 10** Choose **Save** from the **Command** menu or click the **disk icon** on the toolbar to save the inclusion region.




---

**Note** If you made an error in defining the inclusion area, click the area. The selected area is outlined by a dashed aqua line. Next, click the **X icon** on the toolbar. The area is removed from the floor map.

---

- Step 11** Select the **Location Regions** check box if it is not already selected. If you want it to apply to all floor maps, click **Save settings**. Close the Layers configuration page.
- 

## Defining an Exclusion Region on a Floor

To further refine location calculations on a floor, you can define areas that are excluded (exclusion areas) in the calculations. For example, you might want to exclude areas such as an atrium or stairwell within a building. As a rule, exclusion areas are generally defined within the borders of an inclusion area.

To define an exclusion area, follow these steps:

- 
- Step 1** Choose **Monitor > Site Maps**.
  - Step 2** Click the name of the appropriate floor area.

- Step 3** From the **Select a command** drop-down list, choose **Map Editor**.
- Step 4** Click **Go**.
- Step 5** At the map, click the **purple box** on the toolbar.
- Step 6** Click **OK** in the message box that is displayed. A drawing icon appears to outline the exclusion area.
- Step 7** To begin defining the exclusion area, move the drawing icon to the starting point on the map and click once.
- Step 8** Move the drawing icon along the boundary of the area you want to exclude. Click once to start a boundary line and click again to end the boundary line.
- Step 9** Repeat Step 8 until the area is outlined and then double-click the **drawing icon**. The defined exclusion area is shaded in purple when the area is completely defined.
- Step 10** To define additional exclusion regions, repeat Step 5 to Step 9.
- Step 11** When all exclusion areas are defined, choose **Save** from the Command menu or click the **disk icon** on the toolbar to save the exclusion region.



---

**Note** To delete an exclusion area, click the area to be deleted. The selected area is outlined by a dashed purple line. Next, click the **X icon** on the toolbar. The area is removed from the floor map.

---

- Step 12** Select the **Location Regions** check box if it is not already selected, click **Save settings**, and close the Layers configuration page when complete.

## WebGL Requirements

The CMX analytics provides the ability to view the analytic results in both 2D (Open Street Maps) and 3D (WebGL) environments. This provides improved understanding of results on multiple floor paths or when dwell times are calculated throughout a multi-story building. The 3D environment presents the same information as the 2D environment.

WebGL is an advanced feature that provides graphic capabilities. All browsers do not support WebGL on specific hardware. Verify your browser compatibility at:

<http://get.webgl.org/>

If your browser supports WebGL, then on that website you will see a spinning cube.

If your browser does not support WebGL, you must do the following:

- Update to the most current driver for your video card.
- For Google Chrome, follow the instructions given in the Google Chrome support website.
- For Firefox, follow these steps to enable WebGL:
  - In the browser address line, enter **about:config**
  - In the Search text box, enter **webgl** to filter the settings.
  - Double click **webgl.forceenabled**.
  - Make sure that **webgl.disable** is disabled.
- For Safari, follow these steps to enable WebGL:
  - Download the latest build of Safari browser.

- You must enable the Develop menu and enable the WebGL.
- To enable Develop menu, choose **Safari > Preferences**.
- Click the **Advanced** tab.
- Select the **Show Develop** menu in the menu bar check box.
- Choose **Enable WebGL** from the Develop menu.

**Note**

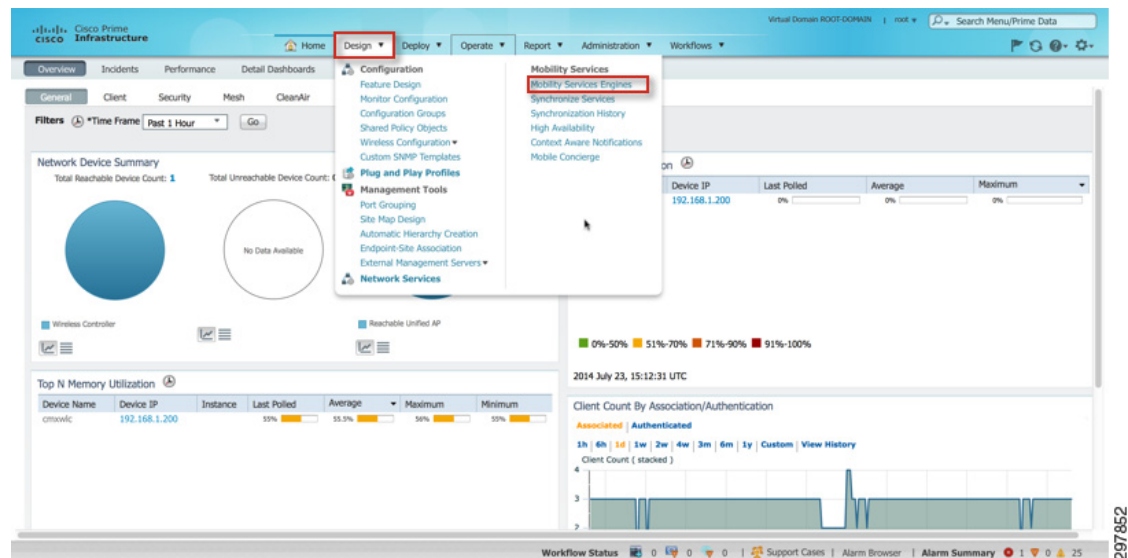
If your system does not support 3D, then the analytic results are displayed only in 2D Open Street Maps view provided that GPS markers are enabled.

Internet Explorer 10 does not have built-in support for WebGL and Microsoft has not announced any plans for implementing it in the future. WebGL support can be manually added to Internet Explorer using third-party plug-ins such as Chrome-frame.

## Adding Mobility Services Engine


**Step 1** Select **Design > Mobility Services Engine**.

**Figure 24-13** Select Mobility Services Engine



**Step 2** From the top right drop down list, select **Add Mobility Services Engine**.

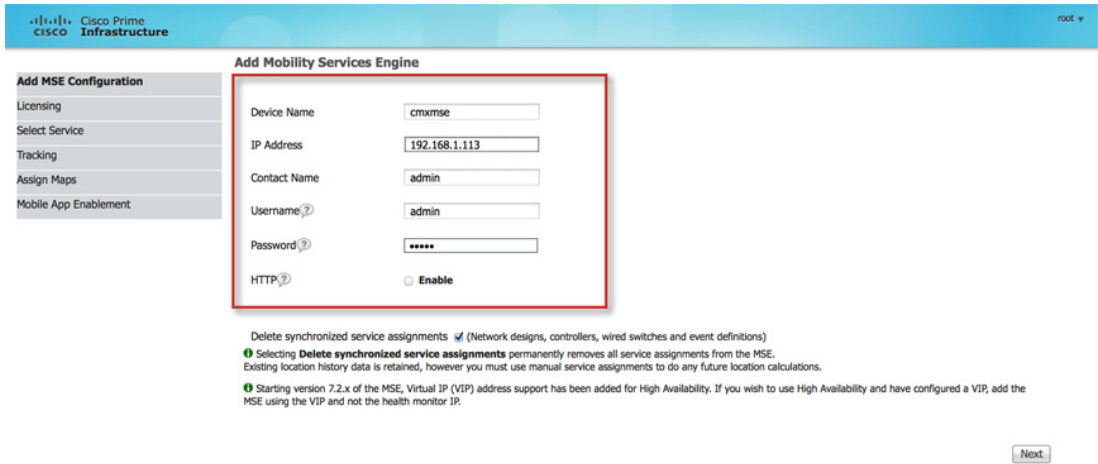
Figure 24-14 Add MSE to PI



297853

**Step 3** Enter the Mobility Services Engine IP Address. Keep the default username and password unless you changed it during MSE setup. In most cases it is the default.

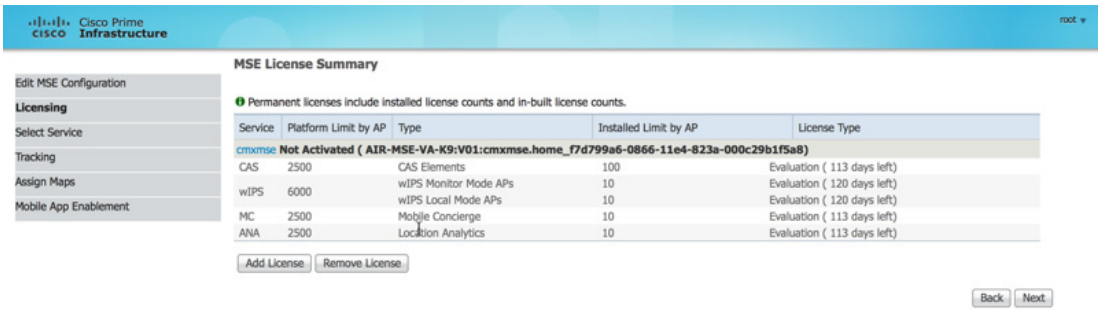
Figure 24-15 Enter MSE Details



297854

**Step 4** The license information on what is available is indicated as shown in Figure 24-16 based on the capacity of the MSE. For the Cisco CMX solution, a valid CMX license is required.

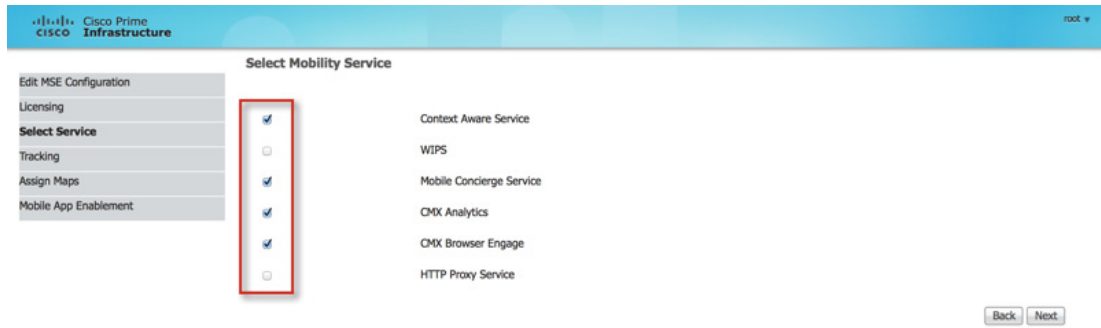
Figure 24-16 Configure Licensing



297855

**Step 5** For the Cisco CMX solution, enable Context Aware Service, CMX Analytics Service, and CMX Browser Engage for Visitor Connect. Optionally, Mobile Concierge service may also be enabled for mobile app communication.

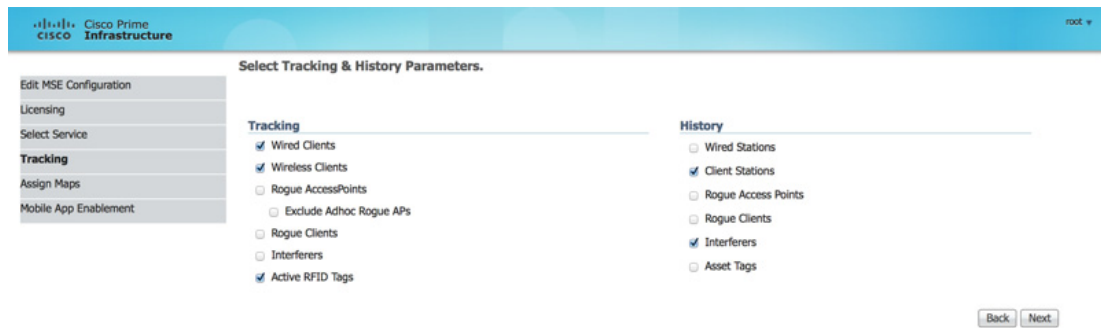
Figure 24-17 Configure Services for MSE



297856

**Step 6** Enable tracking at a minimum of Wireless Clients and maintain the history of Client Stations.

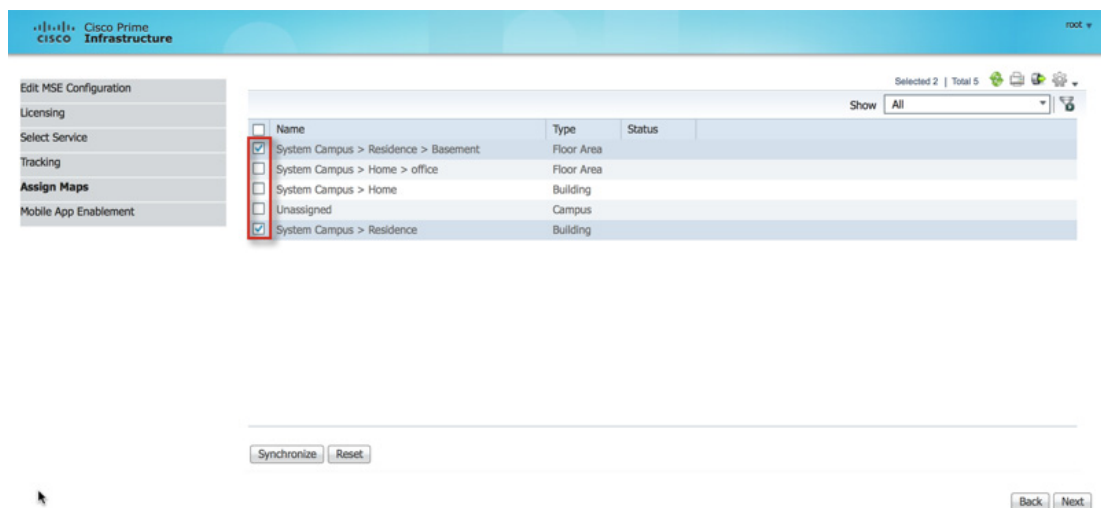
Figure 24-18 Enable Tracking Parameters



297857


**Step 7** Synchronize the relevant maps with the Mobility Services Engine.

Figure 24-19 Synchronize Maps with MSE



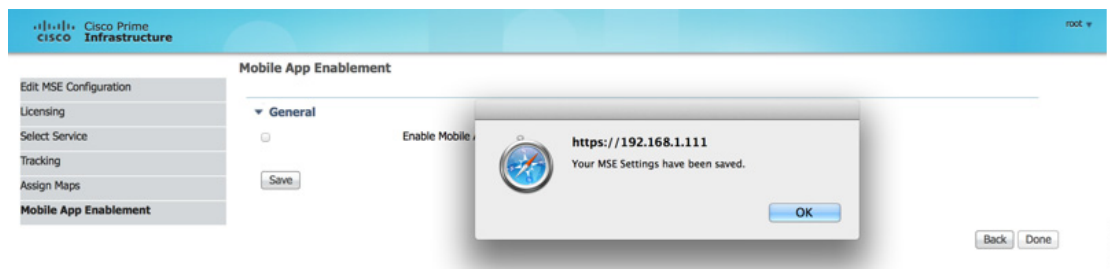
297858

**Step 8** Optionally, enable Mobile App Integration. In this Cisco CMX guide, we are not considering Mobile App integration, so the option has been disabled.

**Figure 24-20** Mobile App Enablement Disabled

297859

**Step 9** Once selected, click **OK** to save the MSE settings.

**Figure 24-21** Save MSE Settings

297860

**Step 10** Verify that the MSE is listed and reachable from the Mobility Services Engine page.

**Figure 24-22** MSE is Reachable via PI

 A screenshot of the Cisco Prime Infrastructure web interface showing the 'Mobility Services Engines' page. The page has a navigation bar with 'Home', 'Design', 'Deploy', 'Operate', 'Report', 'Administration', and 'Workflows'. Below the navigation bar, there is a search bar and a 'Go' button. The main content area is a table with the following data:
 

Device Name	Device Type	IP Address	Version	Reachability Status	Secondary Server	Mobility Service		
Name	Admin Status	Service Status						
cmase	Cisco Mobility Services Engine - Virtual Appliance	192.168.1.113	8.0.1.70	Reachable	N/A (Click here to configure)	Context Aware Service	Enabled	Up
						WIPS	Disabled	Down
						Mobile Concierge Service	Enabled	Up
						CMX Analytics	Enabled	Up
						CMX Browser Engage	Enabled	Up
						HTTP Proxy Service	Disabled	Down

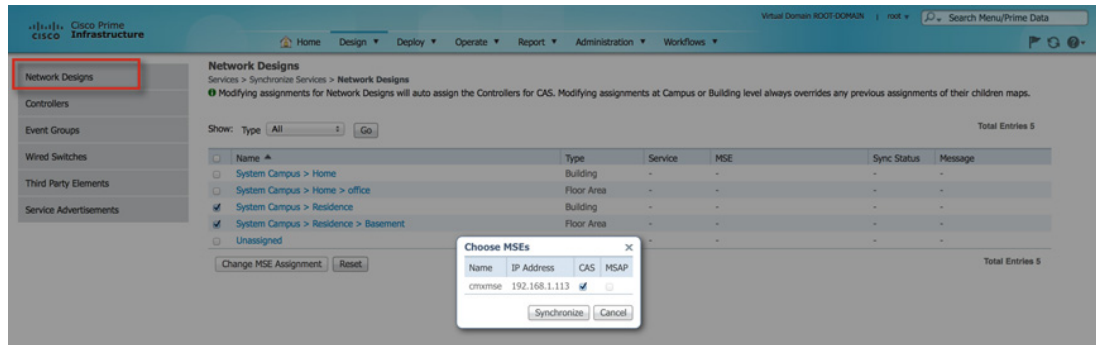
297861

## Synchronizing Controller and Network Designs

**Step 1** Go to **Design > Mobility Sync Services** on Prime Infrastructure.

**Step 2** Select the maps you want to synchronize and click **Change MSE Assignment**. In the pop-up box, select CAS and the MSE to which you want to synchronize the maps and floors.


Figure 24-23 Select Maps to Synchronize



297862

**Step 3** On the left hand side menu, choose **Controllers**.

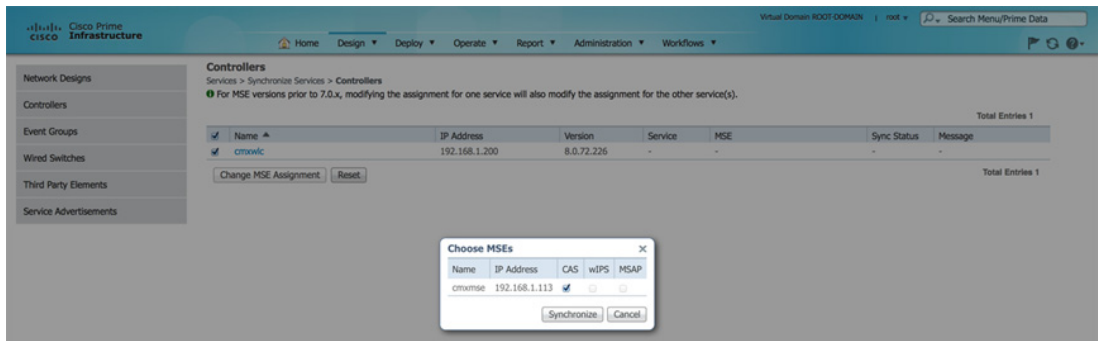
Figure 24-24 Synchronize Controllers with MSE



297863

**Step 4** Select the controllers you want to synchronize with a MSE and click **Change MSE Assignment**. In the pop-up box, select CAS and the MSE to which you want to synchronize the controllers.

Figure 24-25 Synchronize CAS Service with MSE




297864

**Step 5** Click **NSMP status** to ensure that NSMP is up.



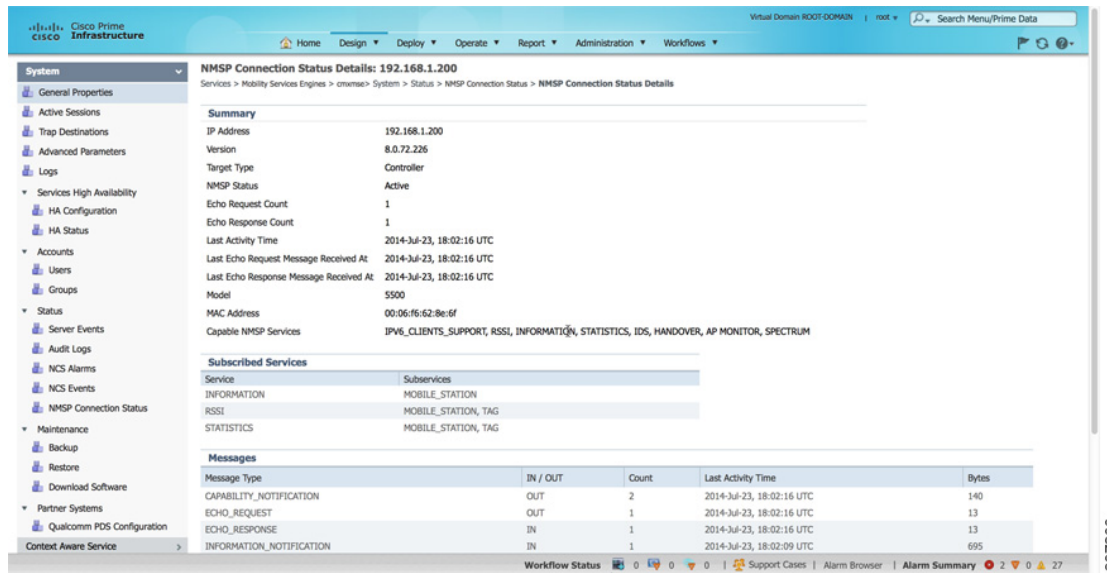
Figure 24-26 Check NMSP Status



Step 6 The NMSP status window should be similar to Figure 24-27. The NMSP status should be listed as Active. If the status is not active, ensure that:

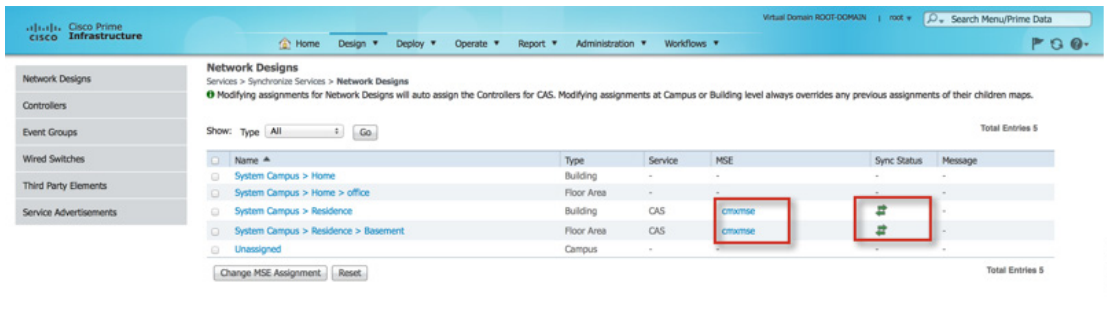
- The MSE and WLC are synchronized to the same time source NTP.
- WLC time is ahead of MSE time.

Figure 24-27 NSMP Status Check on MSE



Step 7 Ensure that the maps are synchronized with the correct MSE.

Figure 24-28 Check Controller Synchronization Status







# Configuring the Mobility Services Engine for CMX


September 4, 2014

Once the Mobility Services Engine (MSE) is configured and synchronized with the WLC, different components of the CMX solution, namely CMX Analytics and CMX Visitor Connect, can be turned on from the MSE UI interface. Note that even if you enabled the services to be turned on as described in [Adding Mobility Services Engine in Chapter 24, “Configuring Cisco Prime Infrastructure,”](#) the services may not be explicitly turned on.

To verify that the services are indeed turned on using the MSE web GUI:

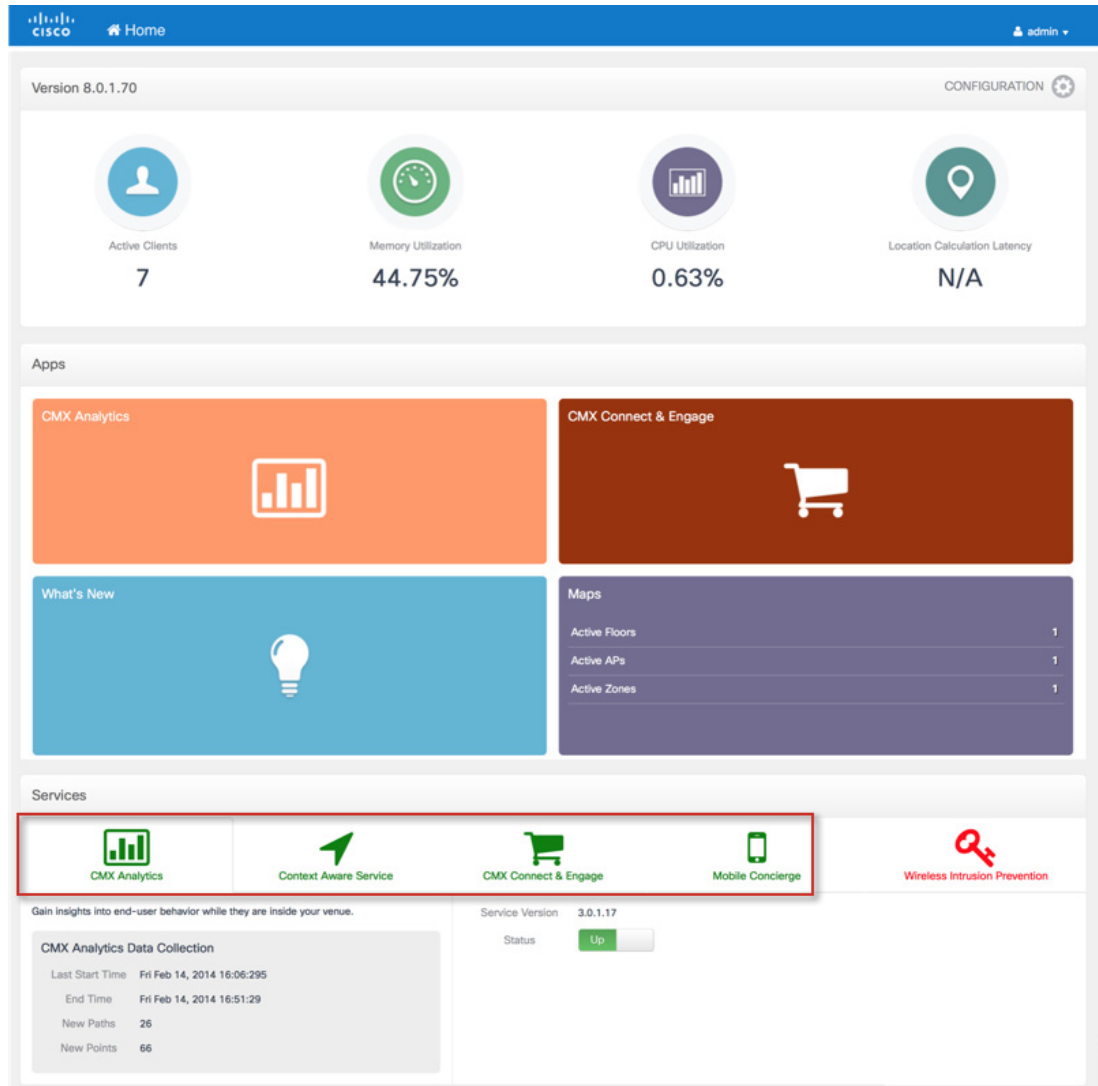
- Step 1** To access the MSE web UI, go to **Design > Mobility Sync Services** within Cisco Prime Infrastructure and click the configured MSE to open the UI in a different page, as shown in [Figure 25-1](#).

**Figure 25-1** Click the MSE URL



- Step 2** Log in to the MSE UI with the username **admin** and the password **admin**. The default username and password can be changed in the MSE UI. On the MSE UI, different services and their status are listed near the end of the page. Select each service and turn them **on** or **off**. For the CMX solution, ensure that CMX Analytics, Context Aware Services, Mobile Concierge Services, and CMX Connect & Engage are turned on. Note that this assumes all of these services are running on a single MSE.

Figure 25-2 MSE Dashboard View

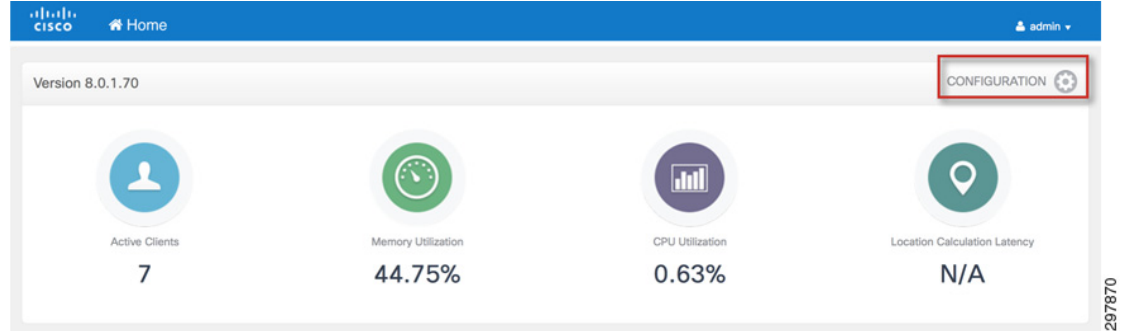


## Verifying CMX Settings

Once CMX has been enabled, it is important to check and verify that the CMX solution components are up and running and that everything is set up properly.

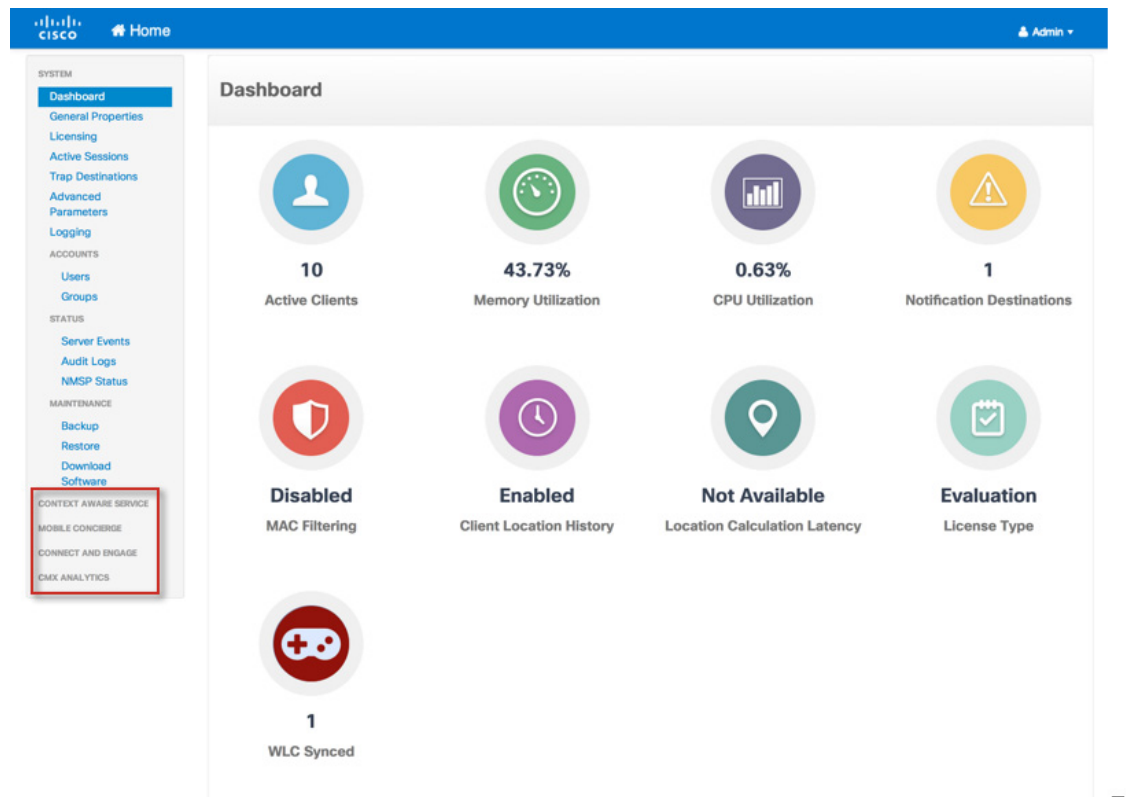
**Step 1** On the MSE UI Dashboard, click the **configuration icon**.

**Figure 25-3** Click the Configuration Icon on the Dashboard



**Step 2** Verify under each of the tabs for Context Aware Service, CMX Analytics, and CMX Connect and Engage that the services are up and are pulling data.

**Figure 25-4** Services List on the CMX Dashboard



**Step 3** Verify that the Context Aware Service is turned on. Under Tracking parameters, ensure that wireless clients are being tracked. Under History parameters, ensure that client history is enabled. CMX Analytics relies on the history of clients being maintained.

Figure 25-5 Figure 36 Ensure Tracking Parameters on CMX

**Tracking Parameters**

Network Location Service Elements Licensed Limit **100**

Elements	Enable Limiting	Limit Value	Active Value	Not Tracked
<input checked="" type="checkbox"/> Wired Clients	<input type="checkbox"/>	0	0	0
<input checked="" type="checkbox"/> Wireless Clients	<input type="checkbox"/>	0	12	0
<input type="checkbox"/> Rogue Access Points	<input type="checkbox"/>	0	0	0
<input type="checkbox"/> Exclude Adhoc Rogue APs	<input type="checkbox"/>			
<input type="checkbox"/> Rogue Clients	<input type="checkbox"/>	0	0	0
<input type="checkbox"/> Interferers	<input type="checkbox"/>	0	0	0
<input checked="" type="checkbox"/> Active RFID Tags	<input type="checkbox"/>	0	0	0

**Save**

Figure 25-6 History Parameters for CMX

**History Params**

Archive for: 30 1 - 365 days

Prune data starting at: 23 hours 50 minutes and also every 1440 minutes

Enable History Logging of Location Transitions for:

- Client Stations
- Wired Stations
- Asset Tags
- Rogue Access Points
- Rogue Clients
- Interferers

**Save** **Cancel**

**Step 4** Verify that under **Connect and Engage > Setup**, the MSE is listed as one of the CAS MSEs. If its not listed, it is important to use the **Add** button and add the MSE IP address. In most cases this should be configured automatically.

Figure 25-7 CAS Service Setup


**Setup**

**Add CAS MSE**

MSE Name	IP address	User Name	Delete
192.168.1.113	192.168.1.113	admin	<input type="button" value="X"/>

- Step 5** Verify that under **CMX Analytics > Setup**, the MSE is listed as one of the CAS MSE. If its not listed, it is important to use the **Add** button and add the MSE IP address. In most cases this should be configured automatically.

**Figure 25-8** *Analytics Services Setup*



297875

## Configuring Role-Based Access Control (RBAC) on the MSE

The MSE itself has its own role-based access control (RBAC) separate from the CMX Connect & Engage service. For role-based access control of the CMX Connect & Engage service, see [Chapter 27, “Configuring RBAC on CMX Connect & Engage.”](#)

To configure RBAC, the MSE administrator must first log in to the MSE via the graphical user interface (GUI).


The following provides an example of the URL to access the MSE Home page.

`https://<MSE_IP_Address>/mseui/apps`

MSE\_IP\_Address is the IP address of the MSE server.

[Figure 25-9](#) shows an example of the login screen which should be displayed.

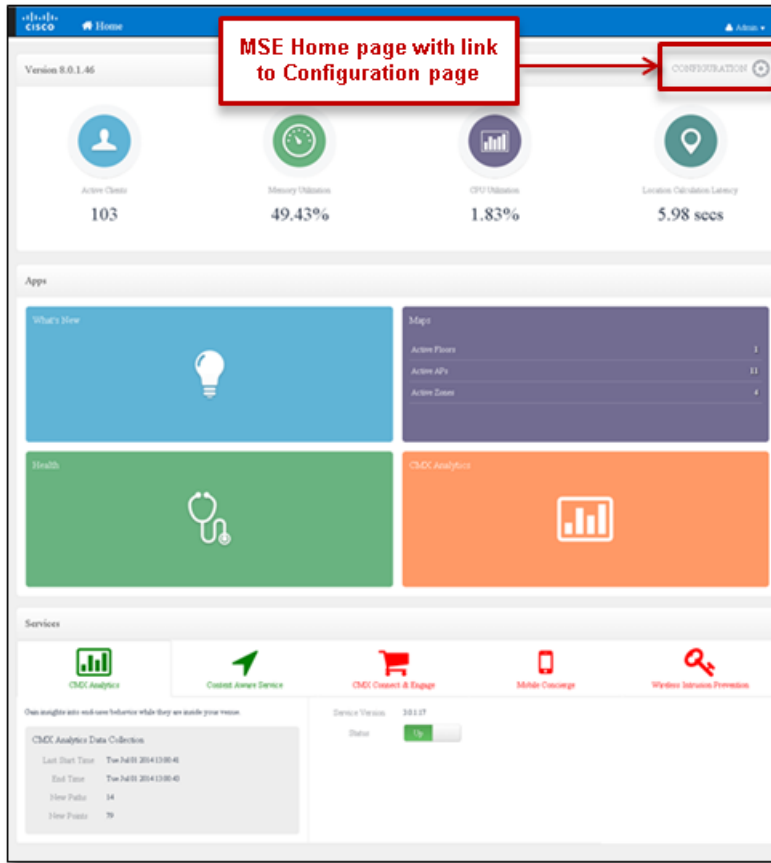
**Figure 25-9** *MSE Login Page*



297876

Upon logging in, the MSE administrator is automatically taken to the MSE Home page, as shown in [Figure 25-10](#).


Figure 25-10 Example of MSE Home Page with Link to Configuration



To configure role-based access control, the MSE administrator needs to click the **Settings icon** in the upper right corner of the page to display the MSE Dashboard page, as shown in [Figure 25-11](#).



Figure 25-11 Example of MSE Dashboard

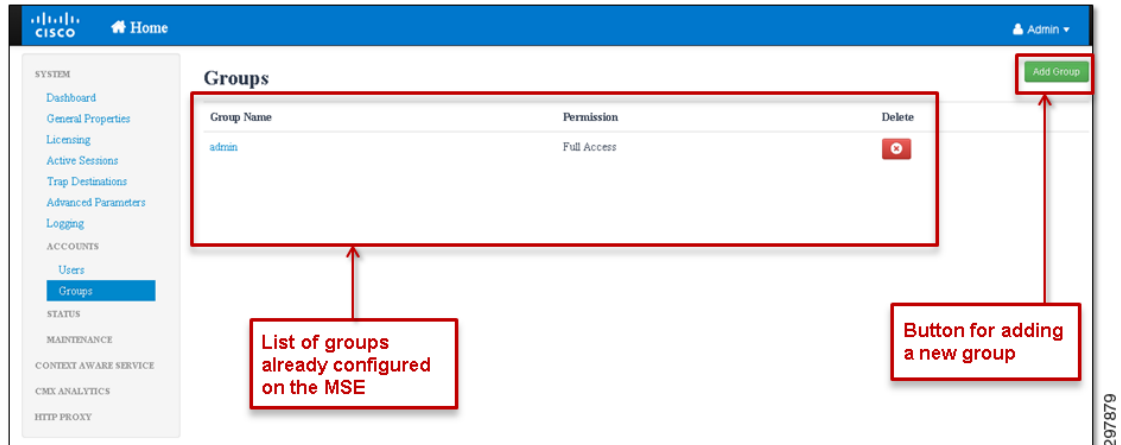


The panel on the left side of the page has four main topics (with sub-topics under several of the main topics) for configuration of the MSE:

- System
- Context Aware Service
- Connect & Engage
- CMX Analytics

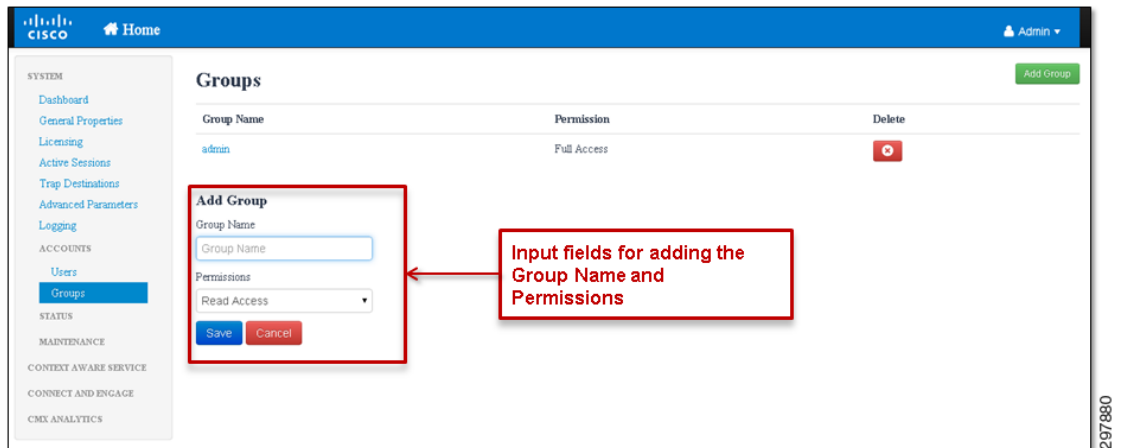
Role-based access control is configured through the Accounts sub-topic under the System topic, as shown in Figure 25-11. The MSE administrator must first configure one or more Groups by clicking the **Groups** link under the Accounts sub-topic located in the panel on the left side of the page, which displays the Groups page, as shown in Figure 25-12.

Figure 25-12 Example of Groups Page



When the MSE administrator clicks the **Add Group** button, the Groups page is modified to include input fields for the Group Name and Permissions, as shown in Figure 25-13.

Figure 25-13 Example of the Groups Page Showing the Add Group Fields



Once the MSE administrator has added the Group Name, they can select the Permissions from the drop down menu:

- **Read Access**—Provides only the ability to view information on the MSE.
- **Write Access**—Provides the ability to make modifications to the configuration of the MSE.
- **Full Access**—Provides complete administrative access, including the ability to enable and disable services, shutdown or reload the MSE, and upgrade MSE code versions.

Clicking the **Save** button adds the new group with desired permissions. Clicking **Cancel** cancels the addition of the new group.

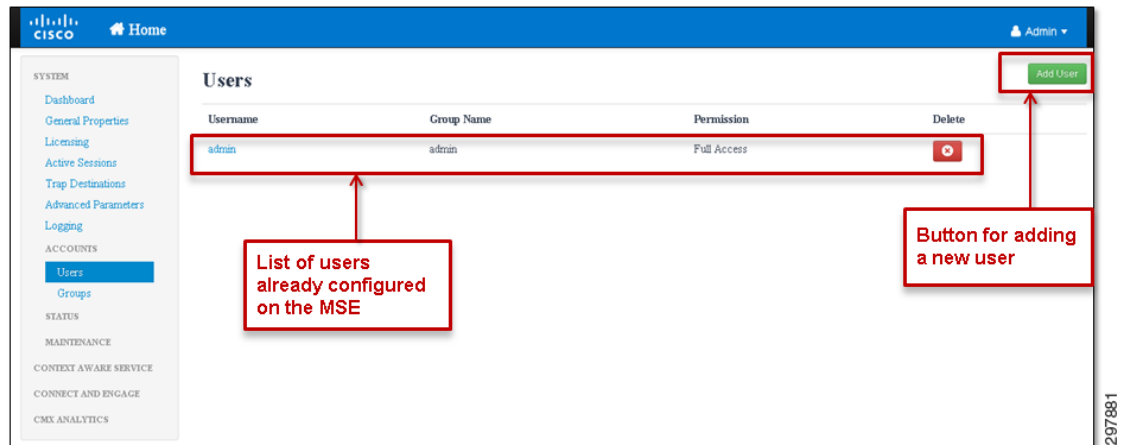
**Note**

As of MSE version 8.0, it is not recommended to utilize Read Access groups on the MSE. CMX Presence Analytics currently does not participate in Role-Based Access Control (RBAC) on the MSE. Hence any userid which is part of a Read Access group is also able to add/modify/delete CMX Presence Analytics configuration. Note that the use of only Write Access and Full Access groups means that any non-IT

personnel who require access to the CMX Analytics dashboard, analytics tab, or reports may also have access to add/modify/delete CMX Analytics (location and presence) configuration. One way to mitigate some of this risk is for IT personnel to download CMX Analytics Reports and email them to non-IT personnel at regular intervals. Alternatively, the list of non-IT personnel—such as store operations managers, marketing executives, etc.—who have direct access to the MSE for CMX Analytics should be kept tightly controlled.

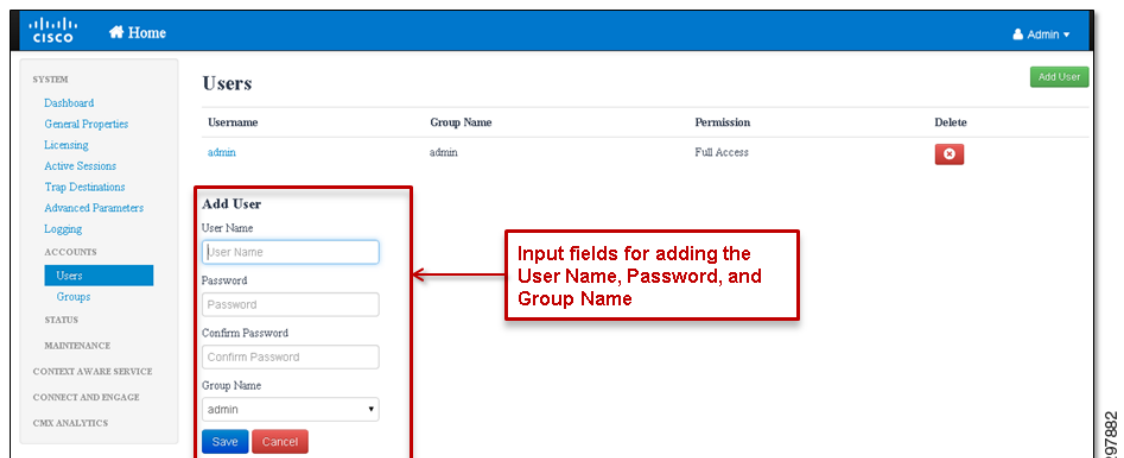
Once the MSE administrator has added the new group, they can add individual user accounts to the group by clicking the **Users** link under the Accounts topic located in the panel on the left side of the page. This displays the Users page, as shown in Figure 25-14.

**Figure 25-14** Example of the Users Page



When the MSE administrator clicks the **Add User** button, the Users page is modified to include input fields for the User Name, Password, and Group Name, as shown in Figure 25-15.

**Figure 25-15** Example of the Users Page Showing the Add User Fields



Once the MSE administrator has added the User Name and Password, they can select the **Group Name** from the drop-down menu. Only groups which were previously configured appear in the drop-down menu.

Clicking the **Save** button adds the new user with desired permissions. Clicking **Cancel** cancels the addition of the new user.



## Configuring CMX Analytics

---

September 4, 2014

This chapter highlights the configuration options available for the three main functional areas of CMX Analytics—Dashboard, Analytics, and Reports. With the CMX Analytics Dashboard, adding, modifying, or deleting pages and widgets are discussed. Each of the different types of analysis, and their configuration options are discussed within CMX Analytics. Finally, each of the different types of reports and their configuration options are discussed with CMX reports.

### Logging In

CMX Analytics is accessed by establishing an HTTPS session directly to the MSE which runs the CMX Analytics service and logging in with a userid and password which has Full Access privileges. Access control to the MSE is discussed in [Configuring Role-Based Access Control \(RBAC\) on the MSE in Chapter 25, “Configuring the Mobility Services Engine for CMX.”](#)

The following provides an example of the URL used to access the CMX Analytics home page.

```
https://<MSE_Name_or_IP_Address>/ui/#login
```

MSE\_Name\_or\_IP\_Address is the name or IP address of the MSE server which runs the CMX Analytics service.

[Figure 26-1](#) shows an example of the login screen which is displayed.

Figure 26-1 CMX Analytics Login Page

Upon logging in, the CMX Analytics administrator is automatically taken to the CMX Analytics Dashboard page, as shown in Figure 26-2.

Figure 26-2 Example of the CMX Analytics Dashboard Page

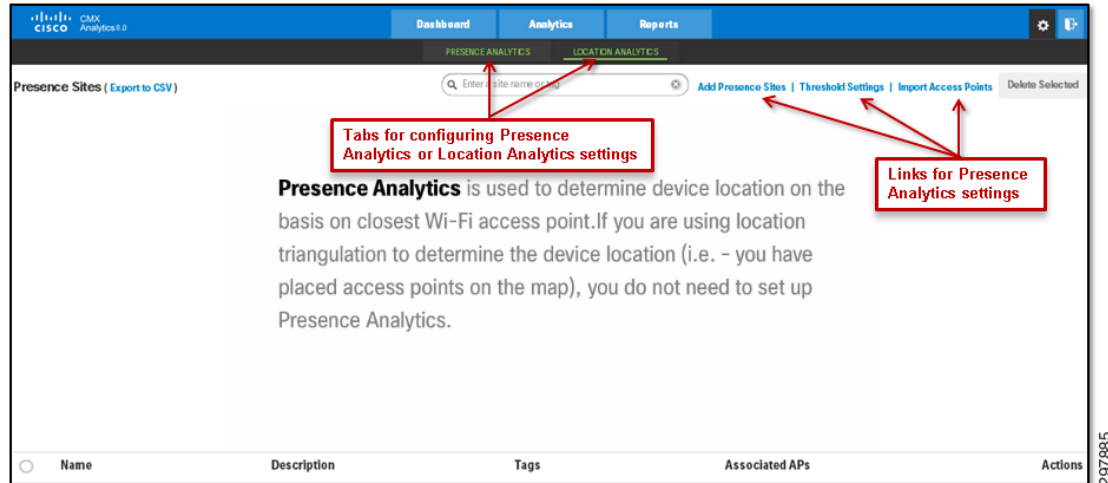


The tabs at the top of the page show the three main functional areas of CMX Analytics—Dashboard, Analysis, and Reporting.

## Configuring CMX Presence Analytics

To configure the settings within Presence Analytics, the administrator needs to click the **Settings icon** in the upper right corner of the screen shown in Figure 26-2, which displays the screen in Figure 26-3.

Figure 26-3 CMX Analytics Settings



The black menu bar across the top of the Settings page has two tabs—Presence Analytics and Location Analytics. By default, Presence Analytics is selected, displaying the configured Presence Sites.

From the Presence Sites page, the CMX Analytics administrator can choose one of the three links located at the top of the page:

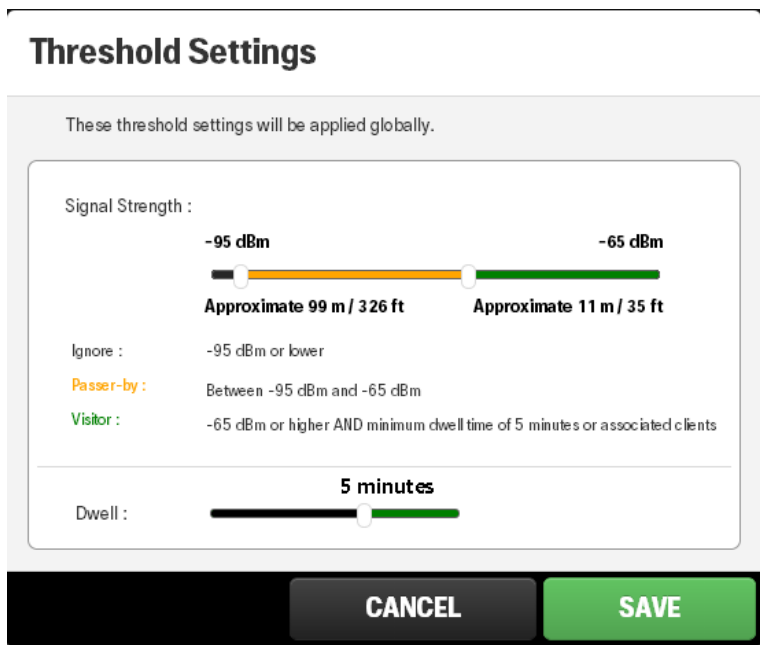
- Add Presence Sites is used to create a new Presence Site, assign access points to the site, and override the Low RSSI Threshold and High RSSI Threshold values for the particular site.
- Threshold Settings is used to set the global default values for the Low RSSI Threshold, High RSSI Threshold, and Minimum Dwell Time for all Presence Sites.
- Import Access Points is used to import access points from Cisco Prime Infrastructure.

Since the Threshold Settings set the global default values for the Low RSSI Threshold, High RSSI Threshold, and Minimum Dwell Time for all Presence Sites, it is discussed first.

## Threshold Settings

From the Presence Sites page, the administrator can click the **Threshold Settings** link to display the Threshold Settings popup window, as shown in [Figure 26-4](#).

Figure 26-4 Example of the Threshold Settings Popup Window



The Threshold Settings popup window includes a horizontal Signal Strength bar with two sliding settings. These settings refer to wireless client signal strength received by the AP or APs associated to the Presence Site.

The left most setting determines the cutoff point between when a wireless client is ignored by Presence Analytics and when the wireless client is considered a passer-by. This setting is also referred to as the Low RSSI Threshold. A wireless client whose signal strength is below the Low RSSI Threshold (as indicated by the black section of the horizontal bar) is ignored. A wireless client whose signal strength is above the Low RSSI Threshold (as indicated by the orange section of the horizontal bar) is considered a passer-by. The Low RSSI Threshold value can be set between -99 dBm by sliding the left most setting all the way to the left and -85 dBm by sliding the left most setting all the way to the right. The default value is -95 dBm.

The right most setting assists, along with the minimum dwell time setting, in determining the cutoff point between when a wireless client is considered a passer-by and when the wireless client is considered a visitor. This setting is also referred to as the High RSSI Threshold. A wireless client whose signal strength is below the High RSSI Threshold (as indicated by the orange section of the horizontal bar) is considered a passer-by. A wireless client whose signal strength is above the High RSSI Threshold (as indicated by the green section of the horizontal bar) for longer than the minimum dwell time is considered a visitor. A wireless client associated to the AP is also considered a visitor. The High RSSI Threshold value can be set between -75 dBm by sliding the left most setting all the way to the left and -40 dBm by sliding the left most setting all the way to the right. The default value is -65 dBm.

The Threshold Settings popup window also includes a horizontal Dwell bar with a single sliding setting. This setting is also referred to as the Dwell Time. The Dwell Time assists Presence Analytics, along with the High RSSI Threshold, in determining when a wireless client is considered a passer-by and when the wireless client is considered a visitor. If a wireless client is detected at a power level above the High RSSI Threshold for a time period greater than the Dwell Time, during the Time Period used to classify the wireless device as a visitor (which is a default of 15 minutes), then the wireless device is classified as a visitor. If a wireless device is detected at a power level above the High RSSI Threshold, for a time period less than the Dwell Time, during the Time Period used to classify the wireless device as a visitor, then the device continues to be considered a passer-by.



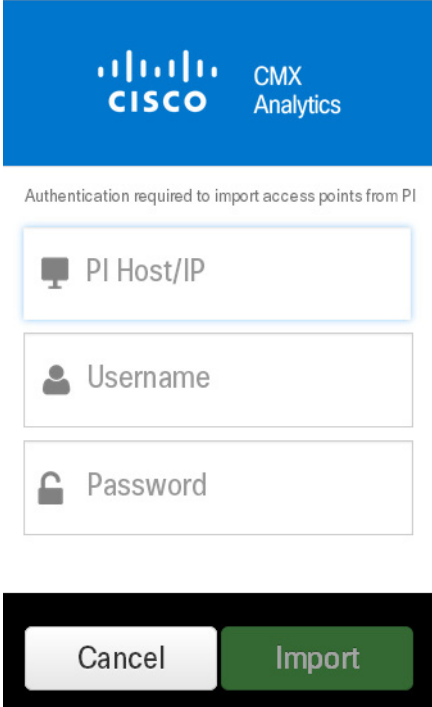
The Dwell Time value can be set between 0 minutes (by sliding the setting all the way to the left) and 8 minutes (by sliding the setting all the way to the right). The default is 5 minutes. The Dwell Time value can only be set globally. It cannot be overridden on a per site basis.

Clicking the **Save** button saves changes to the Threshold Settings and closes the popup window. Clicking **Cancel** cancels the changes and closes the popup window.

## Importing Access Points

By default only APs which appear on maps will be known to the MSE. Since sites which implement Presence Analytics may be small enough such that the CMX Analytics administrator does not desire to create a floor map and import it into Cisco Prime Infrastructure, these APs can be separately imported to the MSE. From the Presence Analytics settings page, the CMX Analytics administrator can click the **Import Access Points** link in order to display the popup window shown in [Figure 26-5](#).

**Figure 26-5** Importing Access Points for CMX Presence Analytics

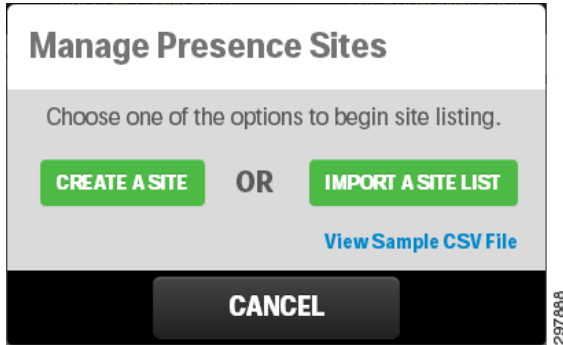


The IP address of the Cisco Prime Infrastructure server, along with a userid and password defined within Cisco Prime Infrastructure, must be configured in the popup window.

## Adding Presence Sites

To add presence sites, the administrator can click the **Add Presence Sites** link shown in [Figure 26-3](#) to display the Manage Presence Sites popup window, as shown in [Figure 26-6](#).

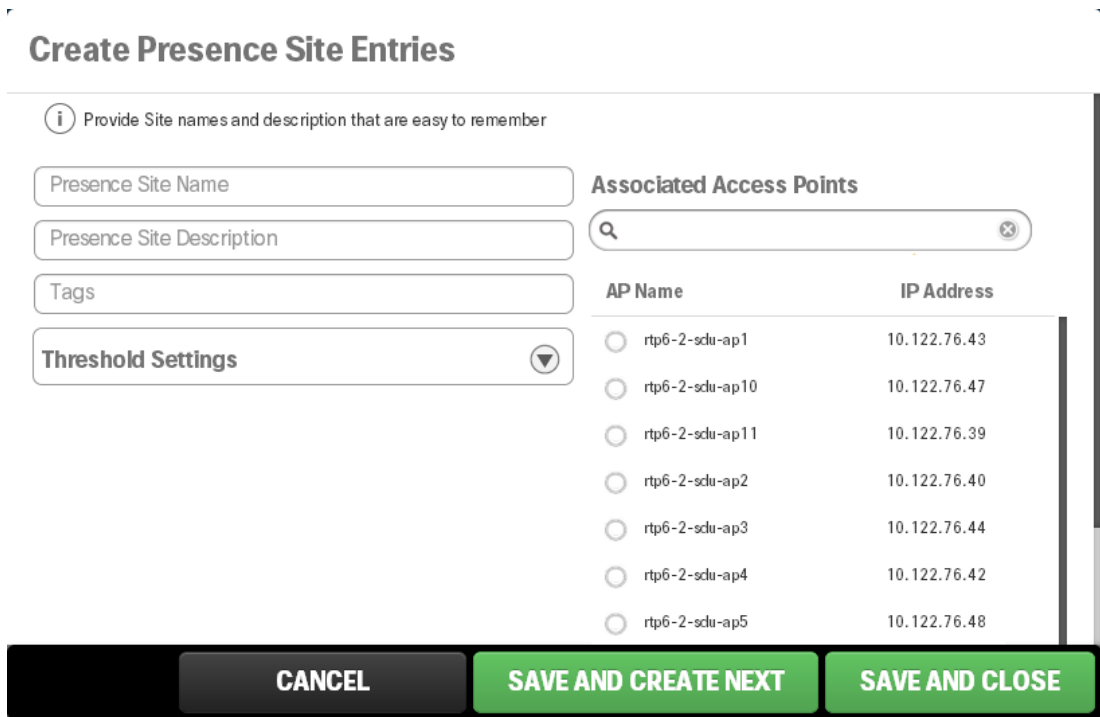
Figure 26-6 Manage Presence Sites Popup Window



The CMX Analytics administrator can manually create presence sites one at a time or import multiple presence sites from a Microsoft Excel spreadsheet.

To create presence sites one at a time, the CMX Analytics administrator can select the **Create A Site** button to display the Create Presence Site Entries popup window, as shown in Figure 26-7.

Figure 26-7 Adding Presence Sites



This popup window includes the following fields:

- Presence Site Name—A unique name that must be selected for each Presence Analytics site.
- Presence Site Description—An optional brief description of the site.
- Tags—An optional list of keywords identifying the site.

Additionally, there is a drop down menu which controls the Threshold Settings which can be customized for the particular Presence Site. The reason why the Threshold Settings can be customized per site is that the RF environment between Presence Sites can be quite different. For example, if one site is a small store surrounded by other stores within a mall, while a second site is a standalone store surrounded by a parking lot, the Threshold Settings may need to be tuned differently.

Finally, the administrator must select the access point or access points which belong to the particular Presence Site from the list of APs which appear in the popup window. Access points which are known to Cisco Prime Infrastructure and the MSE—meaning APs which appear on floor maps or have been imported from Cisco Prime Infrastructure—appear in the popup window. A search field is included in case there are too many APs to be displayed.

Clicking the **Save and Create Next** button creates the new Presence Site and allows the administrator to create another Presence Site through a new popup window. Clicking the **Save and Close** button creates the new Presence Site and returns the administrator to the Presence Analytics settings page. **Cancel** cancels the creation of the new Presence Site.

To create multiple presence sites at once, the CMX Analytics administrator can select the **Import a Site List** button, which guides the administrator to the comma separated value file (.csv format) to upload from their PC. An example of the format of the file is provided by clicking the **View Sample CSV File** link, as shown in [Figure 26-6](#). This downloads a sample file to the administrator's PC. An example of the format of the sample file is shown in [Figure 26-8](#).

**Figure 26-8** Format of Sample File for Creating Multiple Presence Sites

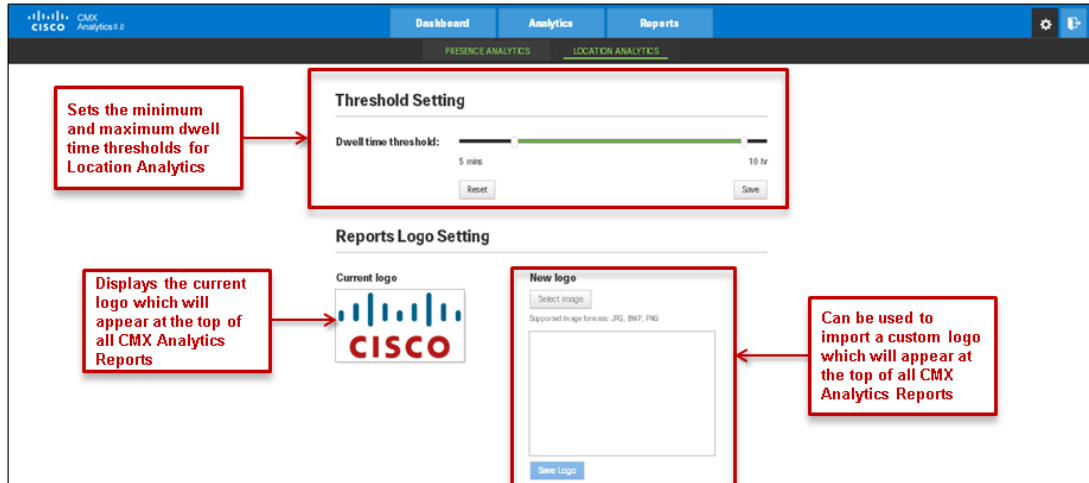
	A	B	C	D	E	F	G	H	I	J	K
1	AP Name	AP IP Address	Base Radio MAC Address	Site Name	Site Description	Tag	RSSI Threshold Low	RSSI Threshold High	Minimum Dwell	Remove	
2	AP 101	192.168.100.101	00:00:00:00:01:01	Store 101	Coffee Shop	Small	-95	-55	5		Site to be added/modified
3	AP 102	192.168.100.102	00:00:00:00:01:02	Store 102	Gift Shop	Medium	-95	-65	5		Site to be added/modified
4	#AP 103	192.168.100.103	00:00:00:00:01:03	Store 103	Restaurant	Large	-95	-75	5		Site commented-out
5	AP 104	192.168.100.104	00:00:00:00:01:04	Store 104	Bakery	Small	-95	-75		3 Y	Site to be removed
6	AP 105	192.168.100.105	00:00:00:00:01:05								AP to be added
7	AP 106	192.168.100.106	00:00:00:00:01:06							Y	AP to be removed
8											
9											

297890

## Configuring CMX Location Analytics

To configure the settings within Location Analytics, the administrator needs to click the **Settings icon** in the upper right corner of the page, as shown in [Figure 26-2](#), which displays the page shown in [Figure 26-3](#). The black menu bar across the top of the Settings page has two tabs—Presence Analytics and Location Analytics. The CMX Administrator must select Location Analytics to display the screen shown in [Figure 26-9](#).

Figure 26-9 CMX Location Analytics Settings Page



The Threshold Setting includes a horizontal bar with two sliding settings. These settings refer to the dwell time of a wireless client seen by Location Analytics.

The Threshold Setting includes a horizontal bar with two sliding settings which are used to filter the data displayed in the output from the Dashboard and Reports tabs. This setting does not apply to the data displayed within output from the Analytics tab.

The leftmost setting determines the minimum cutoff time when a wireless client is filtered out. A wireless client whose dwell time is below the leftmost setting (as indicated by the black section to the left of the sliding setting on the horizontal bar) is not included in the output. A wireless client whose dwell time is above the leftmost setting, but below the rightmost setting (as indicated by the green section of the horizontal bar), is included in the output. The value can be set between 0 minutes by sliding the leftmost setting all the way to the left and 12 hours by sliding the leftmost setting all the way to the right (when the rightmost setting is also slid all the way to the right). The default value is 5 minutes.

The rightmost setting determines the maximum cutoff time when a wireless client is filtered out. A wireless client whose dwell time is above the rightmost setting (as indicated by the black section to the right of the sliding setting on the horizontal bar) is not included in the output. The value can be set between 1 minute by sliding the rightmost setting all the way to the left (when the leftmost setting is also slid all the way to the left) and 24 hours by sliding the rightmost setting all the way to the right. The default value is 10 hours.

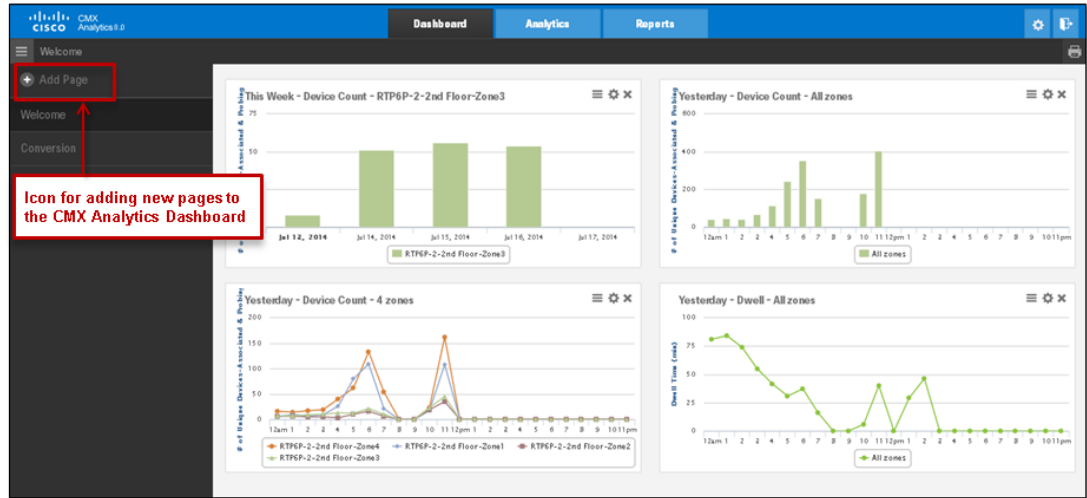
## Configuring the CMX Analytics Dashboard

By default, the CMX Analytics Dashboard tab is selected upon logging into CMX Analytics. Multiple pages, each with multiple widgets, can be configured within the CMX Analytics Dashboard. The name of the currently displayed page appears in the upper left corner of the menu bar at the top of the page, along with an icon for a hidden left panel, which is discussed shortly. Widgets which have been added to the page appear as separate panels within the central part of the page.

### Adding a New Page

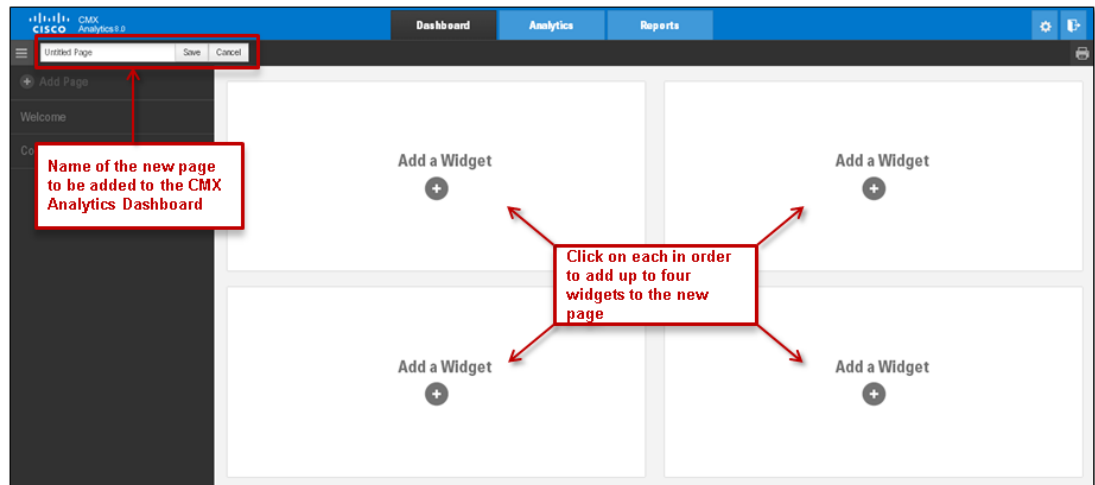
To add a new page, the CMX Analytics administrator must first click the icon in the upper right hand corner of the menu bar across the top of the page, as shown in [Figure 26-2](#), which expands the hidden left panel, as shown in [Figure 26-10](#).

Figure 26-10 Example of MSE Dashboard with Expanded Left Panel



In the left panel, the CMX Analytics administrator can click the **Add Page** icon, which causes the contents of the CMX Dashboard to change, as shown in Figure 26-11.

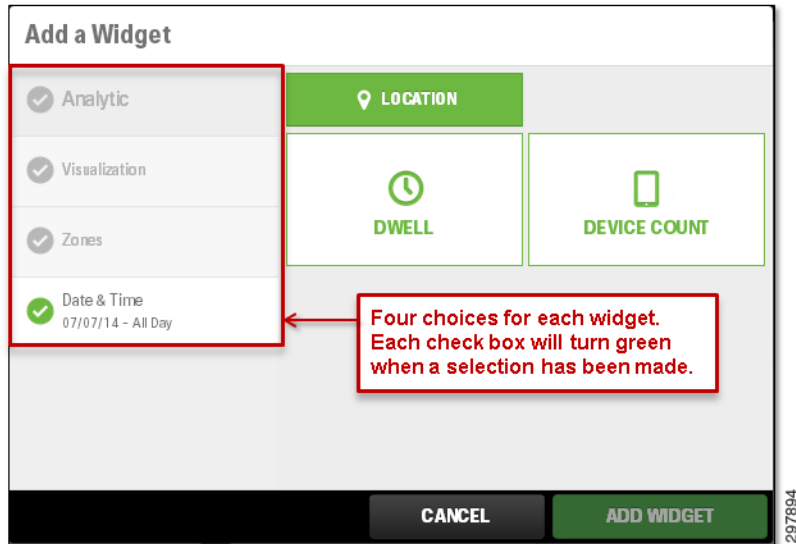
Figure 26-11 Adding a New Page to the CMX Analytics Dashboard



As can be seen in Figure 26-11, a new page with no widgets is displayed. The CMX Analytics administrator should configure a name for the new page. Up to four widgets can be added to each page by clicking one of the four panels in the center of the page.

Clicking one of the **Add a Widget** panels displays a popup screen similar to Figure 26-12.

**Figure 26-12** Adding a Widget to a New Dashboard Page

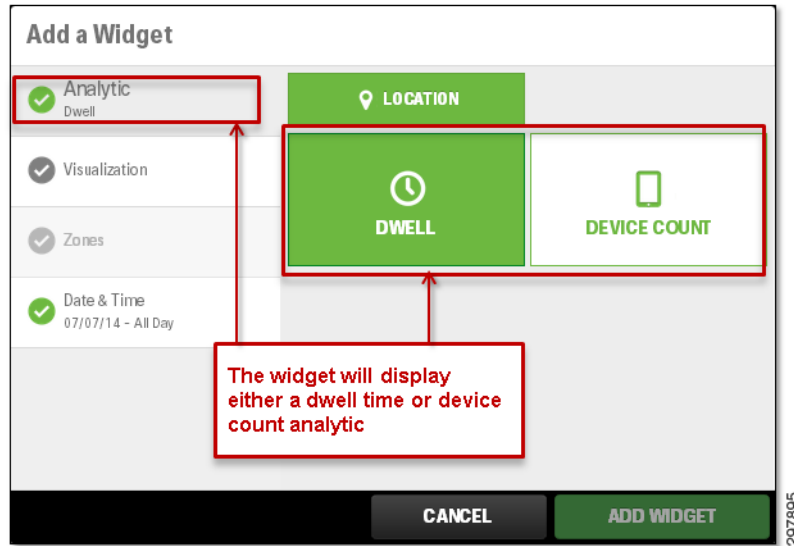


The CMX Analytics administrator needs to make four choices regarding the type of widget to add to the page:

- What type of analytic—Dwell time or device count
- How the analytic is to be visualized—Line chart or bar chart
- Which zone or zones are to be included in the analysis within the widget
- The date and times ranges to be included in the analysis within the widget

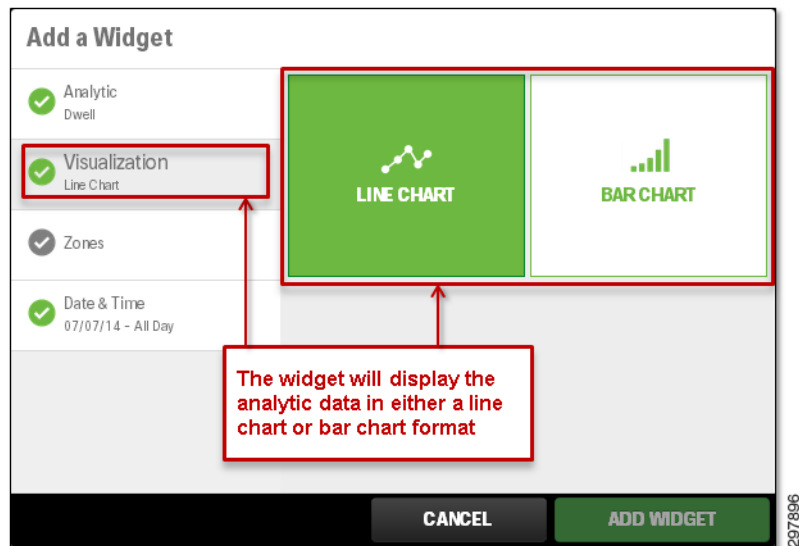
By default the choice begins with the type of Analytic. Clicking either **Dwell** or **Device Count** in [Figure 26-12](#) selects dwell time or device count to be added as the analytic to be covered within the widget. Note that for Presence Analytic sites, a third choice, Conversion Percentage, also appears. Selecting one of the choices causes the Analytic checkbox to turn green, indicating that a selection has been made, as shown in [Figure 26-13](#).

**Figure 26-13** Selection of the Dwell or Device Count Analytic



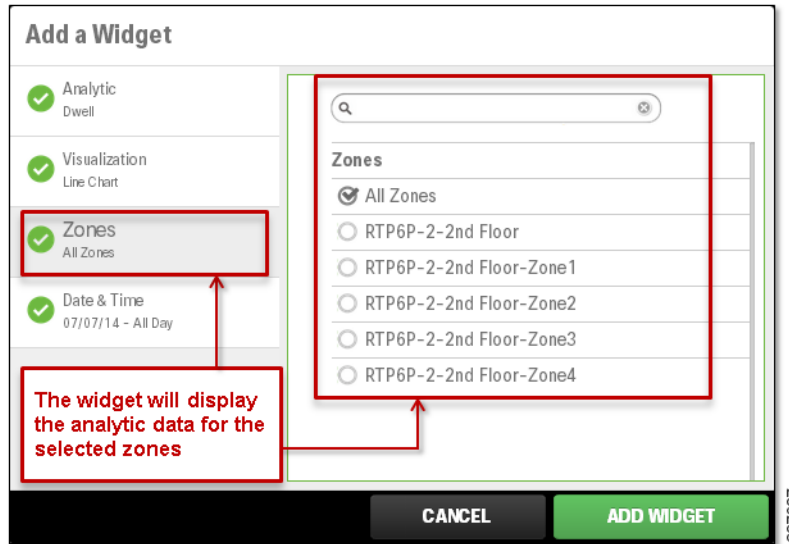
By clicking **Visualization**, the CMX administrator is given a choice to display the analytic data in either a line chart or bar chart format, as shown in [Figure 26-14](#).

**Figure 26-14** Selection of Line or Bar Chart Format for the Widget



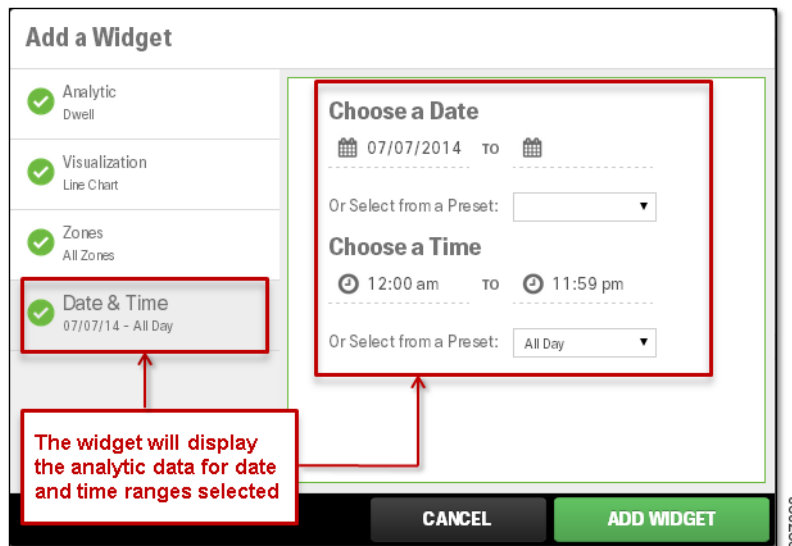
By clicking **Zones**, the CMX administrator is given a choice to include analytic data from a single zone (by selecting the checkbox next to the zone name), multiple zones (by selecting the checkbox next to each desired zone name), or all zones (by selecting the checkbox next to All Zones), as shown in [Figure 26-15](#).

**Figure 26-15** Selection of the Zones from which Analytics Data Will Be Displayed



Finally, by clicking **Date & Time**, the CMX administrator is given a choice regarding what date and time ranges of analytics data are to be included in the widget, as shown in [Figure 26-16](#).

**Figure 26-16** Selection of the Date & Time Ranges of Analytics Data to be Included



The beginning and ending dates can be selected by clicking one of the preset dates: today, yesterday, this week, last week, this month, or last month. Alternatively, the administrator can select the calendar icons under **Choose a Date**. The beginning and ending times can be selected by clicking one of the preset times: All Day, Business Hours, Morning, Lunch Time, Afternoon, or Evening. Alternatively, the administrator can select the clock icons under Choose a Time. By default, the date is set for the current date and the time is set for All Day.



Once the selections for Analytic, Visualization, Zones, and Date & Time have been made (as indicated by a green checkmark next to each), the ADD WIDGET button at the bottom of the popup page becomes green, allowing the CMX Administrator to add the widget to the new page. Note that clicking the **ADD WIDGET** button also automatically adds the new page if this is the first widget added to the page.

## Modifying or Deleting an Existing Page

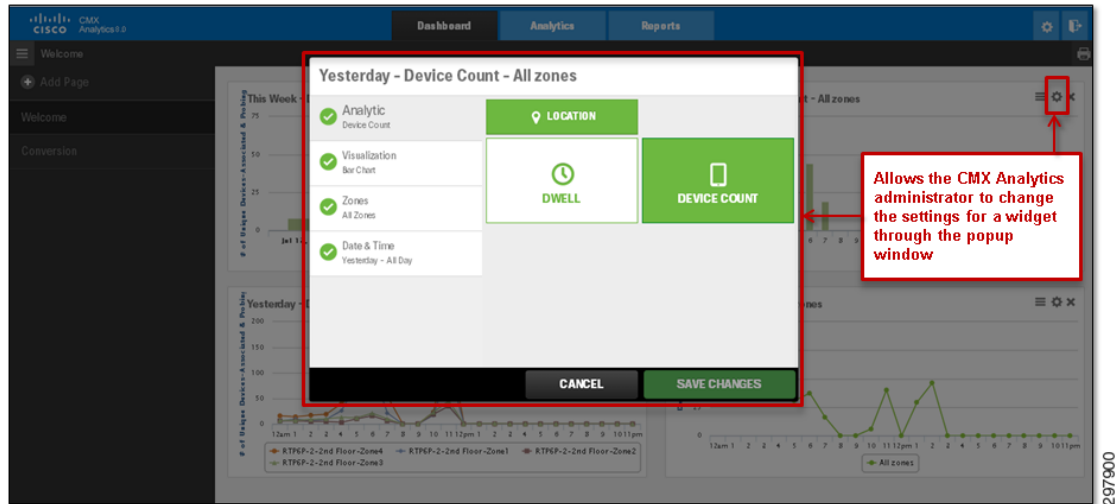
Widgets within an existing page can be modified by clicking the title of the widget or one of the three icons in the upper right corner of each widget. Clicking the title allows the CMX Analytics administrator to customize the title which appears across the top of the widget, as shown in [Figure 26-17](#). Clicking the icon on the left allows the CMX Analytics administrator to either print the widget or export the contents of the widget in .png, .pdf, or .csv file formats.

**Figure 26-17** Customizing the Title and Exporting Contents of a Widget



The settings icon in the middle allows the CMX Analytics administrator to change the settings for the particular widget. Clicking this icon displays the same popup window which was discussed earlier, as shown in [Figure 26-18](#).

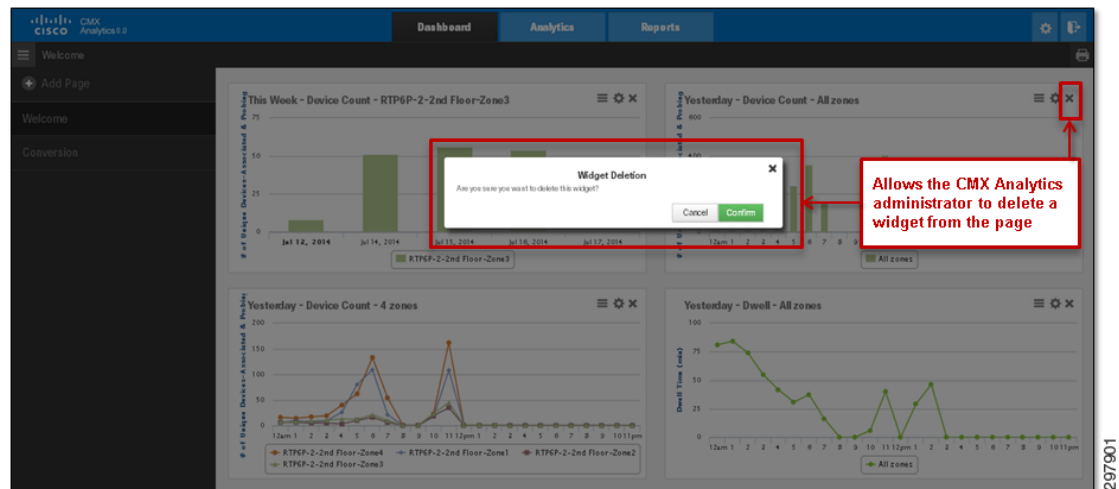
**Figure 26-18** Modifying the Settings of a Widget



The administrator can, for example, change the visualization from a line chart to a bar chart.

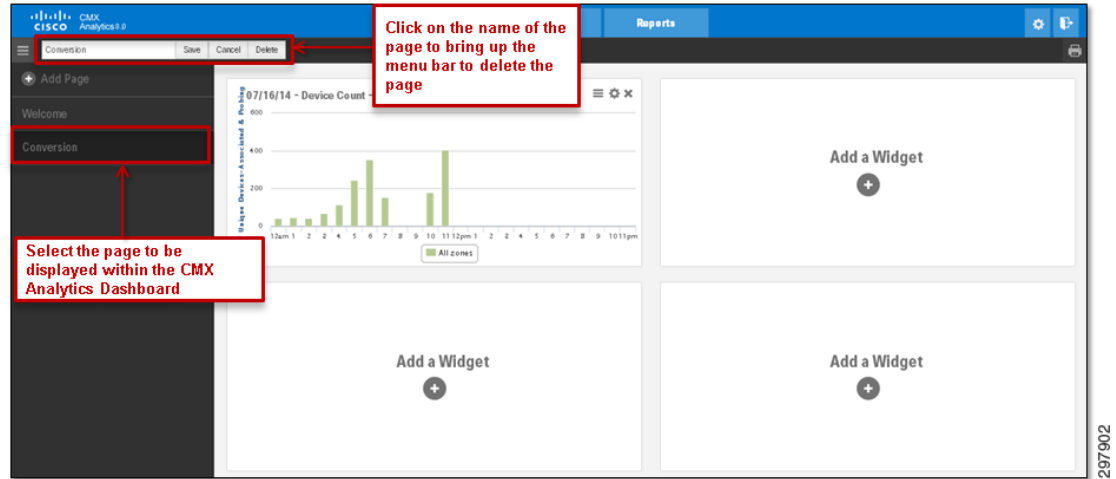
Finally, the delete icon on the right allows the CMX Analytics administrator to delete the particular widget from the page. Clicking this icon displays a popup window which allows the CMX Analytics administrator to confirm the deletion of the widget, as shown in Figure 26-19.

**Figure 26-19** Deleting a Widget from the Page



Once multiple pages have been added to the CMX Analytics Dashboard, the administrator can select which page is currently displayed in the dashboard by clicking the name of the page from the left panel, as shown in Figure 26-20.

**Figure 26-20** Selecting the Page to Be Displayed in the CMX Analytics Dashboard



To delete a page, the CMX Analytics administrator must first select the page to display it. In the menu bar at the top of the page, the administrator must click the name of the page, which displays a menu bar with three choices: Save, Cancel, or Delete. Clicking **Delete** displays a popup window that allows the administrator to confirm the deletion of the page.

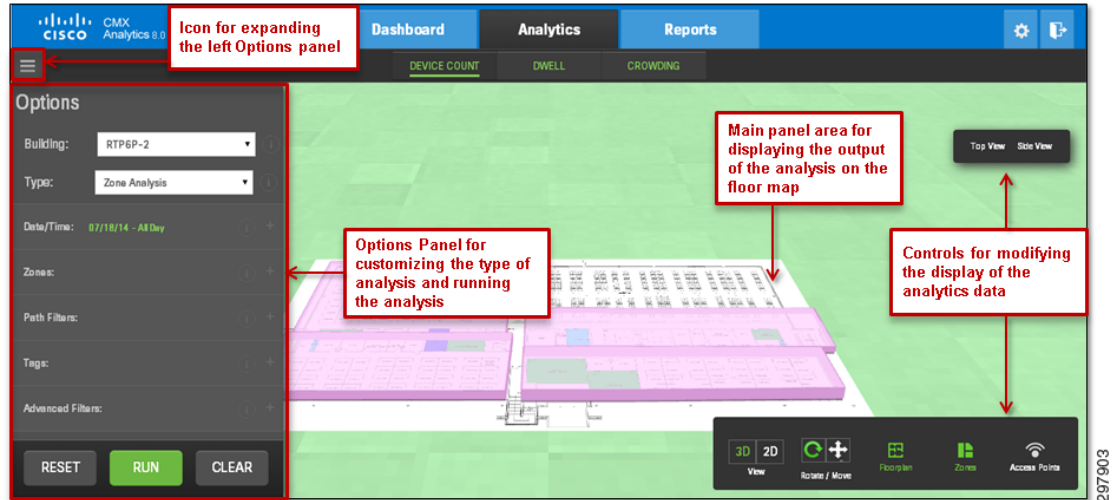
Once the administrator is done configuring the CMX Analytics Dashboard, they can logout by clicking the icon in the upper right corner of the page, as shown in [Figure 26-2](#). This takes the administrator back to the CMX Analytics login popup window.

## Customizing CMX Analysis

As shown in [Figure 26-2](#), the three main functional areas of CMX Analytics are accessed by each of the tabs at the top of the page—Dashboard, Analytics, and Reports. This section briefly discusses the various types of analysis that can be performed by selecting the Analytics tab and how to customize the output from the various types of analysis.

Clicking the **Analytics** tab within CMX Analytics takes the administrator to a screen similar to the one shown in [Figure 26-21](#).

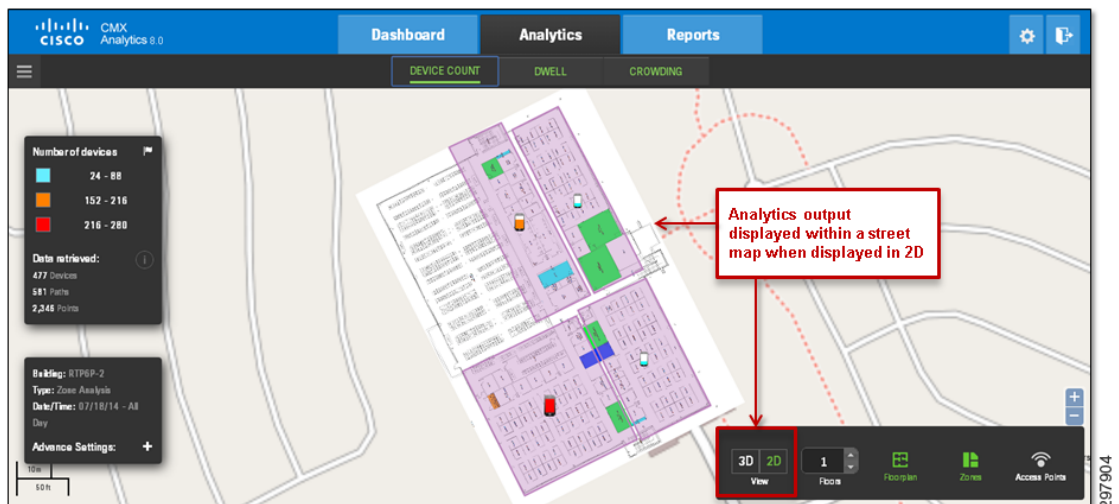
Figure 26-21 Analytics Page



The Analytics page has a hidden left Options panel which can be expanded and contracted via the icon in the upper left corner of the page. The Options panel is used for selecting and customizing the type of analysis to be done and running the analysis. The main panel area is used to display the output of the analysis on the floor map and zones which are chosen within the Options panel. The specific analytic displayed after running the analysis, Device Count, Dwell, or Crowding, can be selected from the three tabs in the black horizontal menu bar at the top of the page. Finally, there are sets of controls which are used to modify how the analytics are displayed. For instance, the floor can be displayed in 3-dimensions (3D) or 2-dimensions (2D) and the floor plan, zones, and access points can be included or not included in the display along with the analytics data.

To display the output in 2D the floor plan imported to CMX Analytics must include GPS coordinates, which is configured through Cisco Prime Infrastructure. Analytics output in 2D automatically displays the building placed on a street map, as shown in Figure 26-22.

Figure 26-22 Example of Analytics Output in 2D

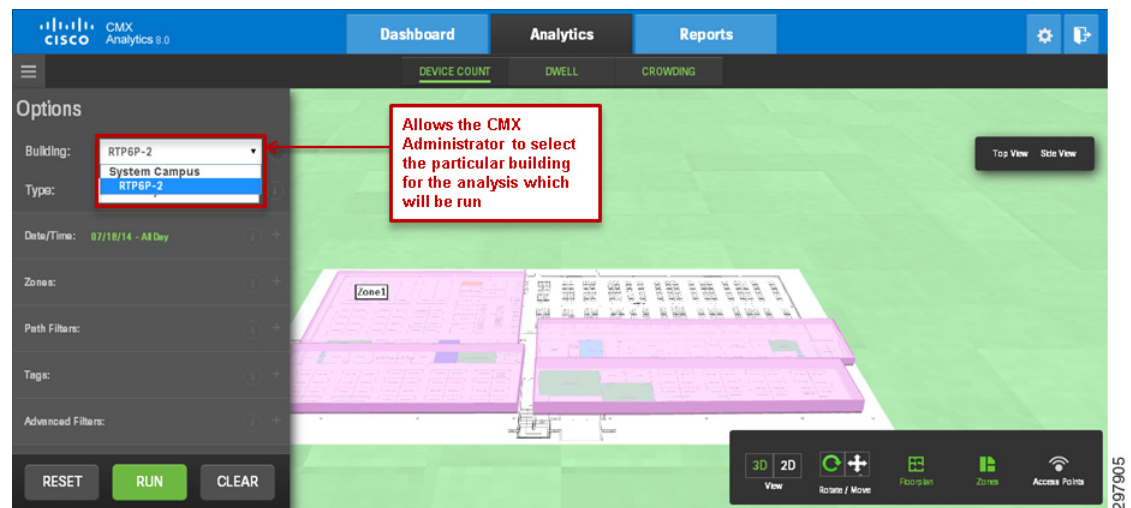


The Options panel has seven selections which are used to customize the analysis to be run before clicking the green **Run** button:

- Building:
- Type:
- Date/Time:
- Zones:
- Path Filters:
- Tags:
- Advanced Filters:

A single MSE can be used to collect analytics data for multiple sites. Different sites appear as separate buildings, floors, and/or zones within CMX Analytics. By clicking the drop down menu next to Building within the Options panel, the CMX Analytics administrator can select the building and floor for the analysis which will be run, as shown in [Figure 26-23](#).

**Figure 26-23** Selection of the Building for Analysis

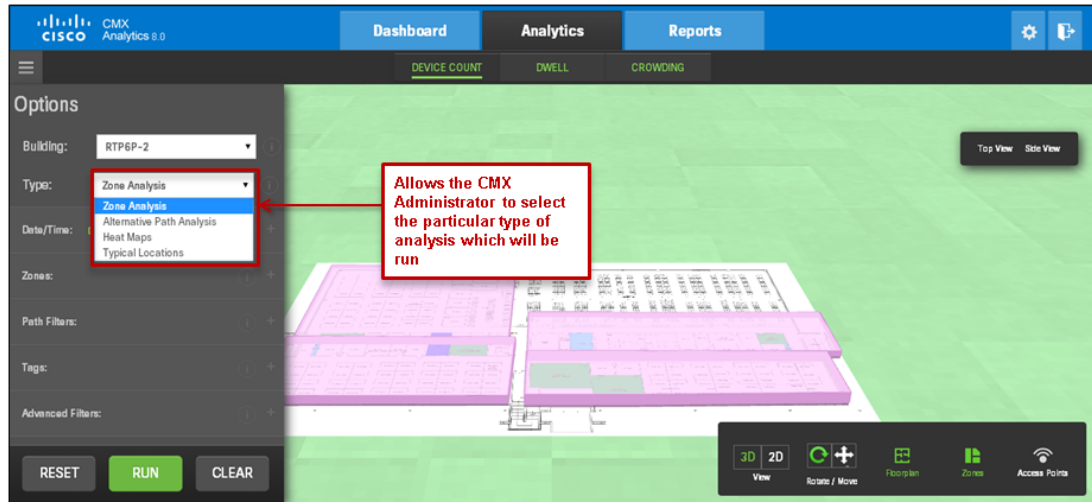


By clicking the drop down menu next to Type: within the Options panel, the CMX Analytics administrator can select from one of the following four types of analysis which can be performed:

- Zone Analysis
- Alternative Path Analysis
- Heat Maps
- Typical Locations

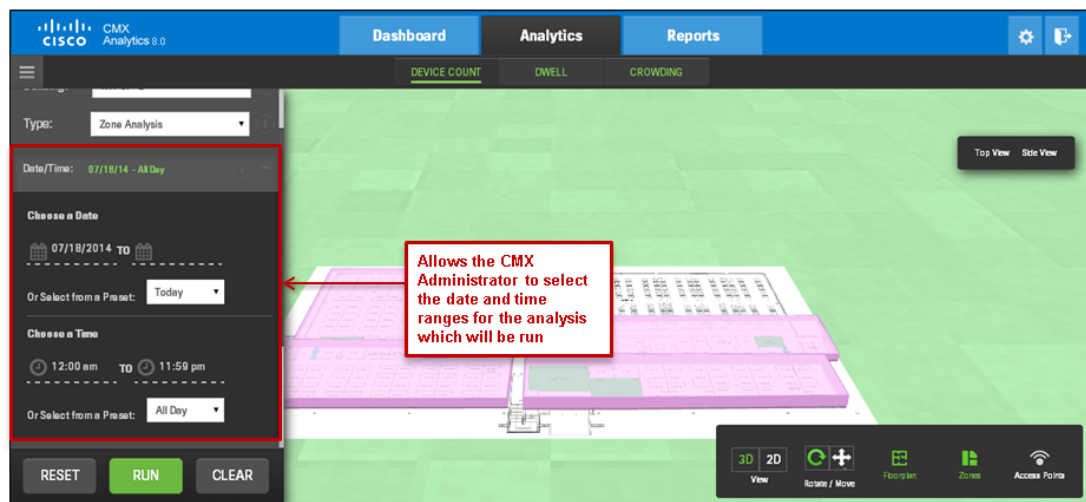
An example is shown in [Figure 26-24](#).

Figure 26-24 Selection of the Type of Analysis



Each of the different types of analysis are discussed in more detail in later sections of this document. By clicking **Date/Time:**, the CMX Analytics administrator is given a choice regarding what date and time ranges of analytics data are to be included in the analysis, as shown in Figure 26-25.

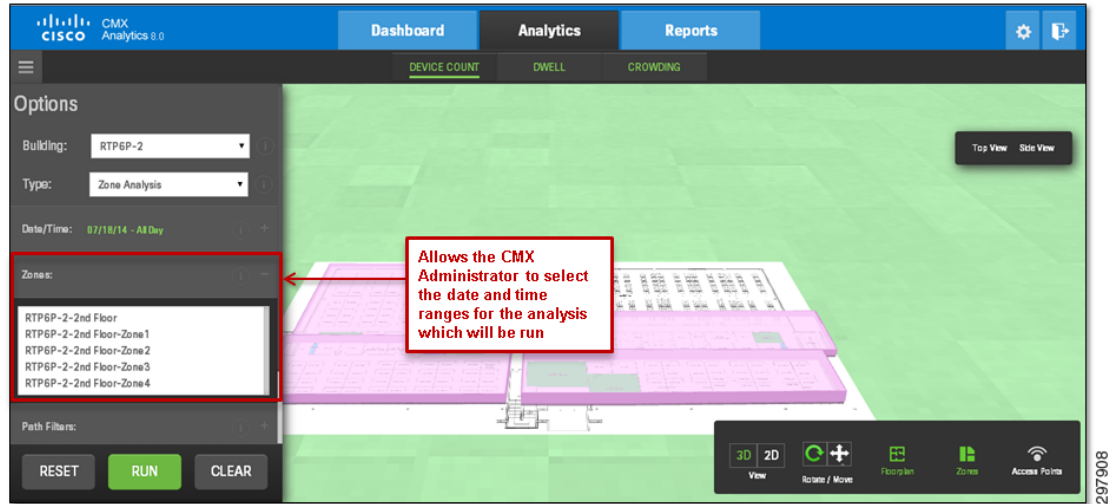
Figure 26-25 Selection of the Date and Time Ranges of the Analysis



The beginning and ending dates can be selected by clicking one of the preset dates: today, yesterday, this week, last week, this month, or last month. Alternatively, the administrator can select the **calendar icons** under Choose a Date. The beginning and ending times can be selected by clicking one of the preset times: All Day, Business Hours, Morning, Lunch Time, Afternoon, or Evening. Alternatively, the administrator can select the **clock icons** under Choose a Time. By default, the date is set for the current date and the time is set for All Day.

By clicking the drop down menu next to Zones: within the Options panel, the CMX Analytics administrator can select the zone or zones for the analysis which will be run. The available zones will correspond to those within the building and floor selected, as shown in Figure 26-26.

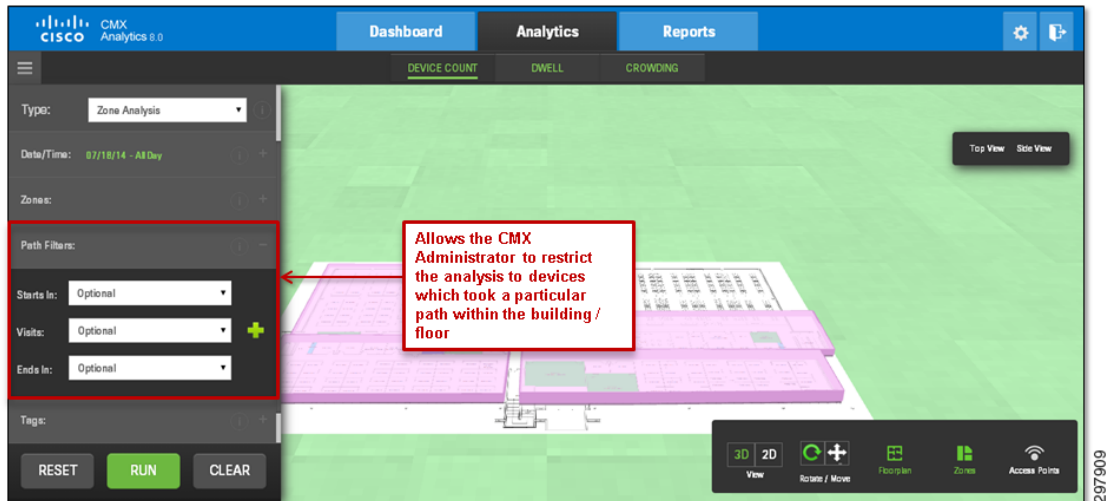
Figure 26-26 Selection of the Zones for the Analysis



CMX Analytics organizes data around the concept of points and paths. Points represent specific locations where a device was calculated to have been detected (via RSSI information within Probe Requests or the use of the FastLocate feature) while it was within the venue. Paths represent the movement of the device between those points. CMX analytics operates on a separate database which is built on data extracted from the MSE location database. Analytics data for a particular device is calculated only after until the path of the device through the particular venue is complete, which is typically an hour after the device is no longer detected within the venue.

Path Filters: allows the CMX Analytics administrator to restrict the analysis to devices which follow a particular path within the building/floor which has been selected. The path can include a beginning zone, one or more intermediate zones, and an ending zone. The keyword, **Optional**, can be selected within each to indicate any zone, as shown in Figure 26-27.

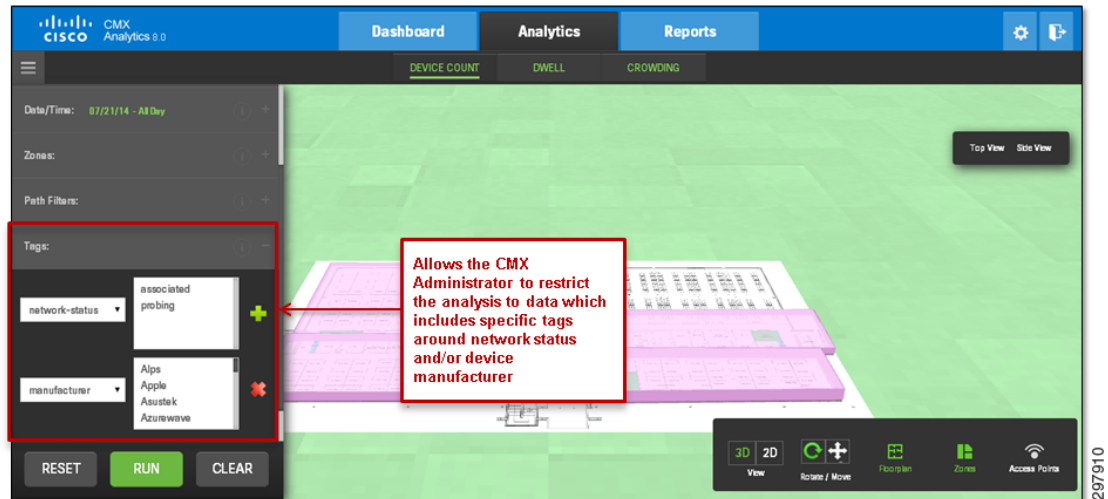
Figure 26-27 Selection of Path Filters for the Analysis





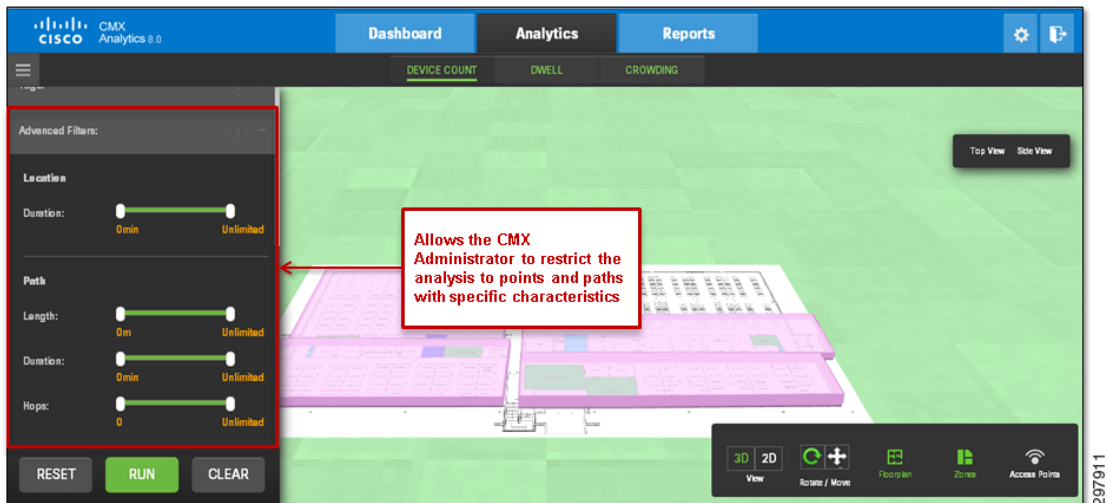
Tags: allows the CMX Analytics administrator to restrict the analysis to data which includes specific tags. The tags which can be configured include network status and device manufacturer. The choices within network status include whether the device is associated with an access point or whether the device is simply sending Probe Requests. The choices within manufacturer include a drop down list of various manufacturers, as shown in [Figure 26-28](#).

**Figure 26-28** Selection of Tags for the Analysis



Finally, by selecting Advanced Filters: within the Options panel, the CMX Analytics administrator can select restrict the analysis to points and paths with specific characteristics, as shown in [Figure 26-29](#).

**Figure 26-29** Selection of Advanced Filters for the Analysis



For example, the sliding bar setting next to Duration: in the figure above can be used to restrict the analysis to those data points which have a minimum and maximum dwell time. Note that this type of dwell time is only measured when a device is stationary, not moving. The sliding bar settings next to Length:, Duration:, and Hops:, in [Figure 26-29](#) can be used to restrict the analysis to those paths which have a specific length in term of meters, duration in terms of minutes or hours, and hops which are specific data points along the path.



The following sections discuss the each of the different types of analysis which can be run.

## Zone Analysis

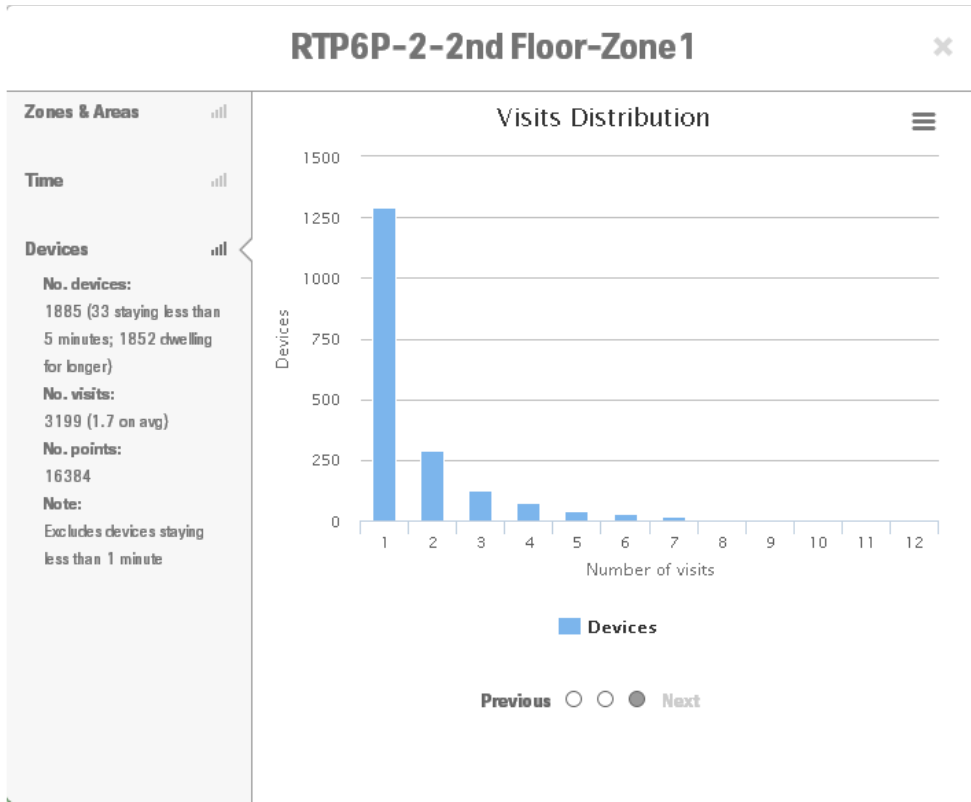
Zone Analysis provides parameters such as dwell time, number of devices, and crowding for each zone selected. [Figure 26-30](#) provides an example of the output from Zone Analysis when device count has been selected as the output of the analysis.

**Figure 26-30** Zone Analysis with Device Count Selected



As can be seen, the total count of devices seen per zone selected, over the time period selected, is displayed. Clicking one of the nodes displays a popup window containing more detailed information, as shown in [Figure 26-31](#).

Figure 26-31 Detailed Output from Zone Analysis—Device Count



The additional information includes a graph showing the distribution of visits of devices seen over the time period which the Zone Analysis was run. For example, [Figure 26-31](#) displays the distribution of visits for all devices seen in one zone over a period of a week. By running this report weekly or monthly, the CMX administrator can see if the number of repeat devices is increasing or decreasing over time. This roughly corresponds to whether or not the number of repeat visitors is increasing or decreasing over time.


[Figure 26-32](#) provides an example of the output from Zone Analysis when dwell time has been selected as the output of the analysis.

Figure 26-32 Zone Analysis with Dwell Time Selected



As can be seen, the median dwell time of devices seen per zone selected, over the time period selected, is displayed. Clicking one of the nodes displays a popup window containing more detailed information, as shown in Figure 26-33.

Figure 26-33 Detailed Output from Zone Analysis—Dwell Time



The additional information includes a graph showing the distribution of dwell times of devices seen over the time period which the Zone Analysis was run. For example, Figure 26-33 displays the distribution of dwell times for all devices seen in one zone over a period of a week. By running this report weekly

or monthly, the CMX administrator can see if the dwell times of devices is increasing or decreasing over time. This roughly corresponds to whether or not visitors are staying for longer or shorter durations within the zone over time.

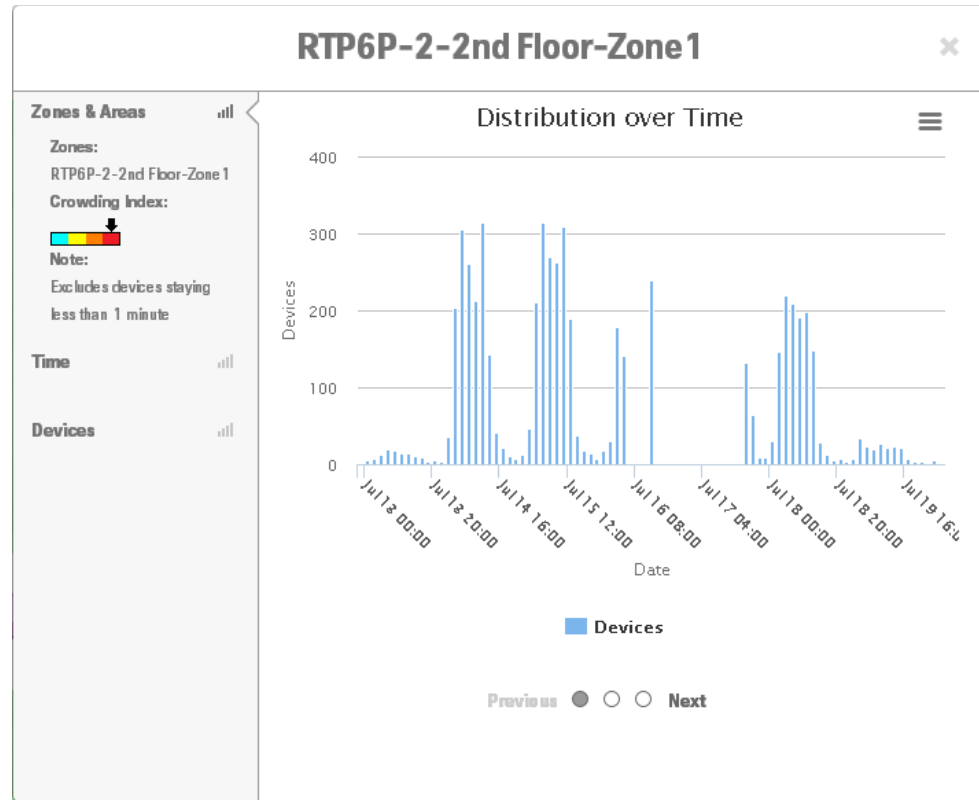
Figure 26-34 provides an example of the output from Zone Analysis when crowding has been selected as the output of the analysis.

**Figure 26-34** Zone Analysis with Crowding Selected



As can be seen, the average crowding factor of devices seen per zone selected, over the time period selected, is displayed. Again, clicking one of the nodes displays a popup window containing more detailed information, as shown in Figure 26-35.

Figure 26-35 Detailed Output from Zone Analysis—Crowding

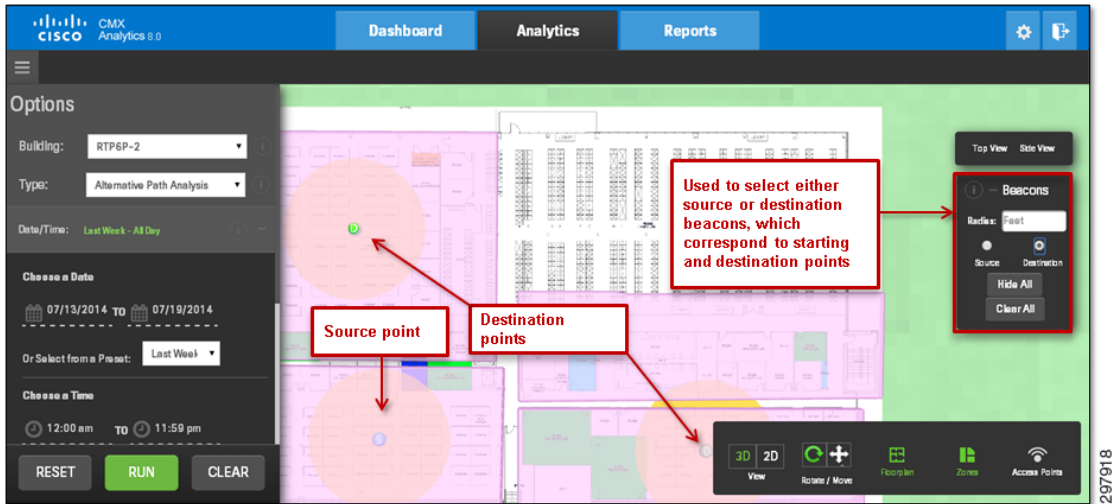


The additional information includes a graph showing a detailed distribution of devices seen over the time period which the Zone Analysis was run. For example, [Figure 26-35](#) displays the distribution of devices seen in one zone broken out by hours and days of a week. By running this report weekly, the CMX administrator can see patterns in peak crowding of visitors from week to week and see if the crowding is increasing or decreasing over time. This could be an indication that additional staffing may be needed to handle the increased crowding of visitors over time during peak hours.

## Alternative Paths Analysis

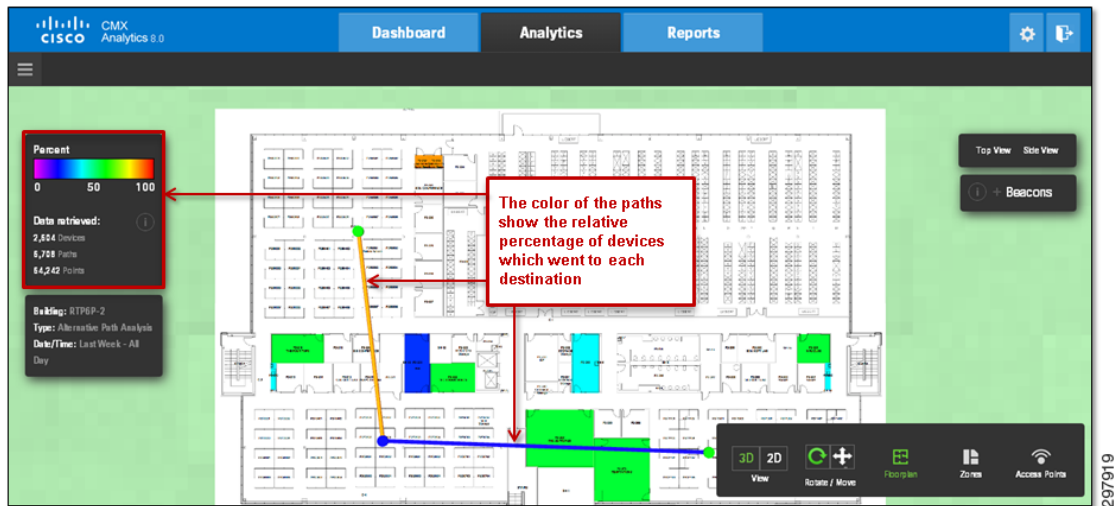
Alternative Paths Analysis shows a breakdown of the percentage of devices going to each destination from each starting point. To run Alternative Paths Analysis, the CMX administrator must first configure one or more source beacons (starting points) and one or more destination beacons (destination points). This is done by selecting either Source or Destination from the panel on the right side of the page and clicking the floor map in the location where the beacon is desired, as shown in [Figure 26-36](#).

Figure 26-36 Alternative Paths Analysis Configuration Example



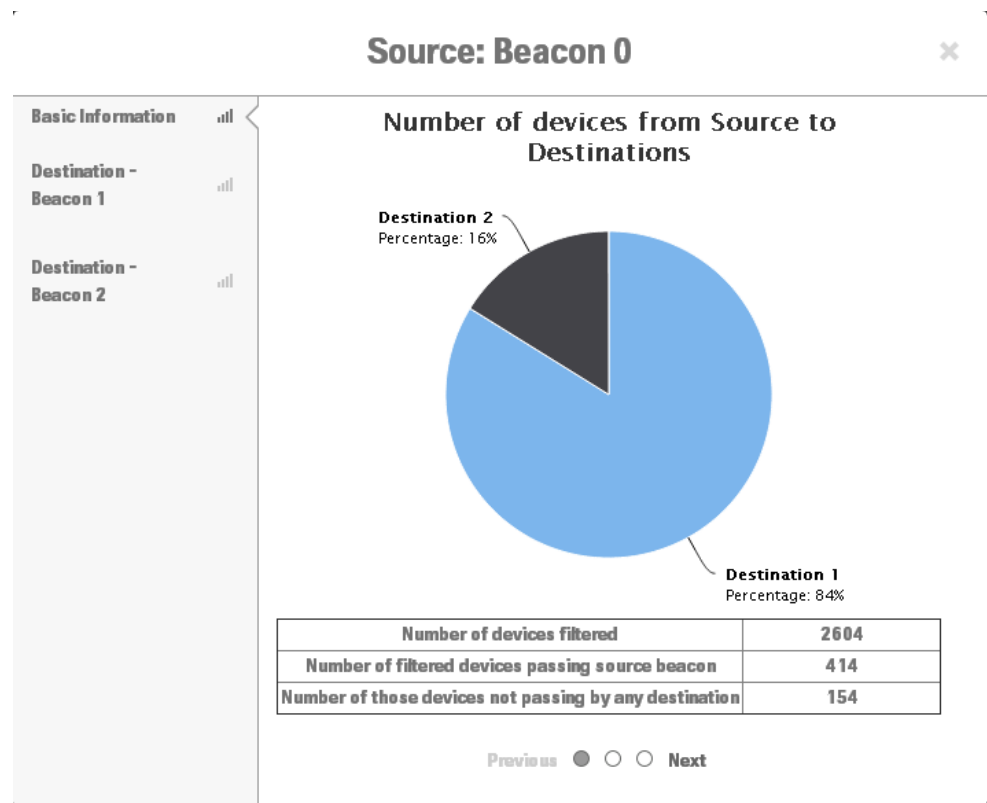
Devices which are first detected within the radius of a source beacon and subsequently detected within the radius of a destination beacon are included within the output of the analysis. An example of the output from Alternative Paths Analysis is shown in Figure 26-37.

Figure 26-37 Alternative Paths Analysis Output Example



As can be seen, the color of the path indicates the relative percentage of devices that were first seen at the source, which then went to each destination. Note that for clarity purposes, Zones are not highlighted in the output shown. Clicking the source beacon displays a popup window containing more detailed information, as shown in Figure 26-38.

Figure 26-38 Detailed Output from Alternative Paths Analysis

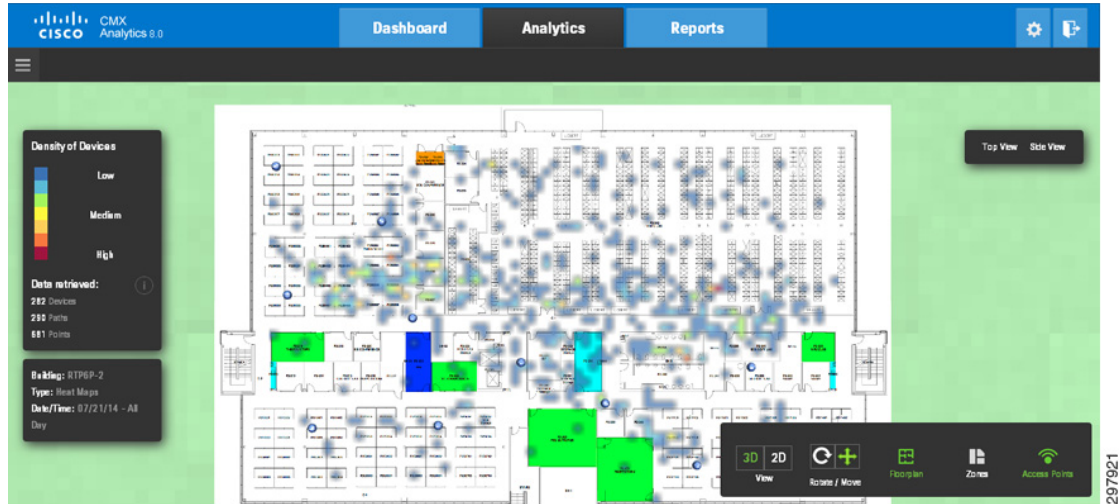


The additional information includes a pie chart showing the specific percentage of devices which then went to each destination. Additionally, the chart shows the total number of devices included in the analysis, the number of those devices which were then detected within the radius of the source beacon, and finally the total number of those devices which did not pass either destination beacons. This type of analysis can be used to determine the movement of visitors to different locations within the venue as they enter the venue. For example, it may be used to determine the percentage of customers which go directly to particular departments upon entering a store.

## Heat Maps

Heat Maps provides a graphical representation of point data which can be viewed on the map in such a way that areas of higher concentration appear darker, as shown in [Figure 26-39](#).

**Figure 26-39** Heat Maps Example Output



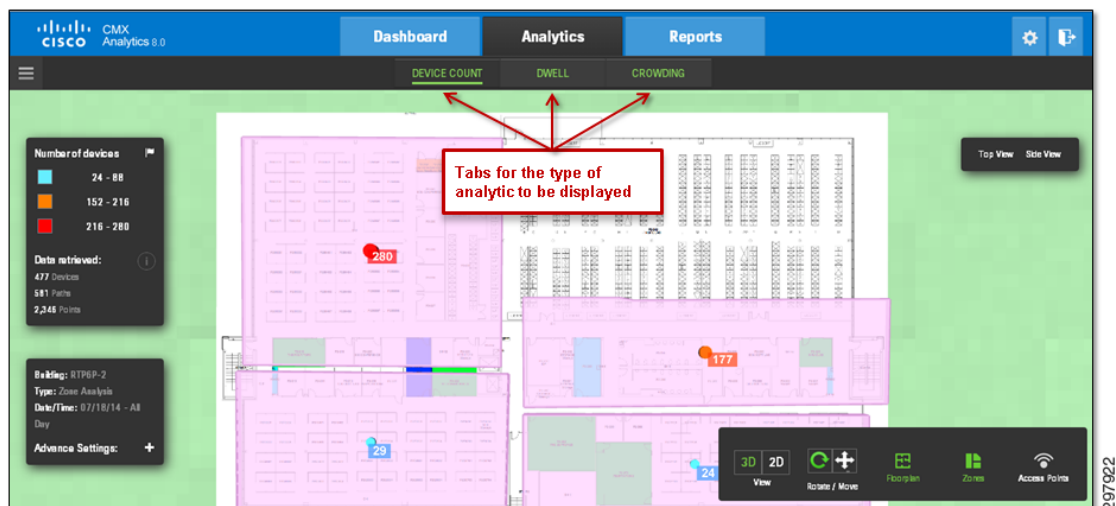
Heat maps can be used to show the density of visitors within various parts of the venue.

## Typical Locations

Typical Locations provides parameters such as dwell time, number of devices, and crowding for different areas of the building. Typical Locations analysis is similar to Zone Analysis. However the analytics information is presented as multiple nodes spread across multiple areas within a particular floor or zone with Typical Locations analysis versus a single node spread across the entire zone with Zone Analysis. These areas are dynamically determined based upon clustering of devices.

Figure 26-40 provides an example of the output from Typical Locations when device count has been selected as the output of the analysis and a single zone is selected.

**Figure 26-40** Typical Locations with Device Count Selected





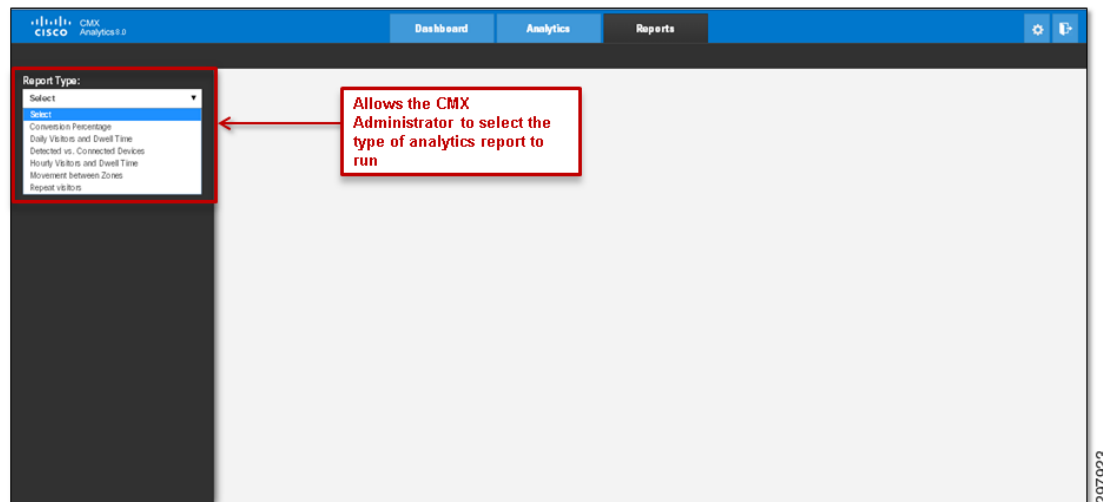
As with Zone Analysis, the output of the analysis can display device count, dwell time, or crowding by selecting one of the three options in the black menu bar across the top of the page. Likewise, clicking one of the nodes displays a popup window containing the same detailed information as was discussed in [Zone Analysis](#).

## Customizing CMX Reports

As shown in [Figure 26-2](#), the three main functional areas of CMX Analytics are accessed by each of the tabs at the top of the page—Dashboard, Analytics, and Reports. This section briefly discusses the various types of analytics reports that can be run by selecting the **Reports** tab and how to customize the output from the various types of reports.

Clicking the **Reports** tab within CMX Analytics displays a screen similar to that shown in [Figure 26-41](#).

**Figure 26-41** CMX Analytics Reports Tab



The panel on the left side of the screen includes a drop down menu allowing the administrator to select from one of the following reports:

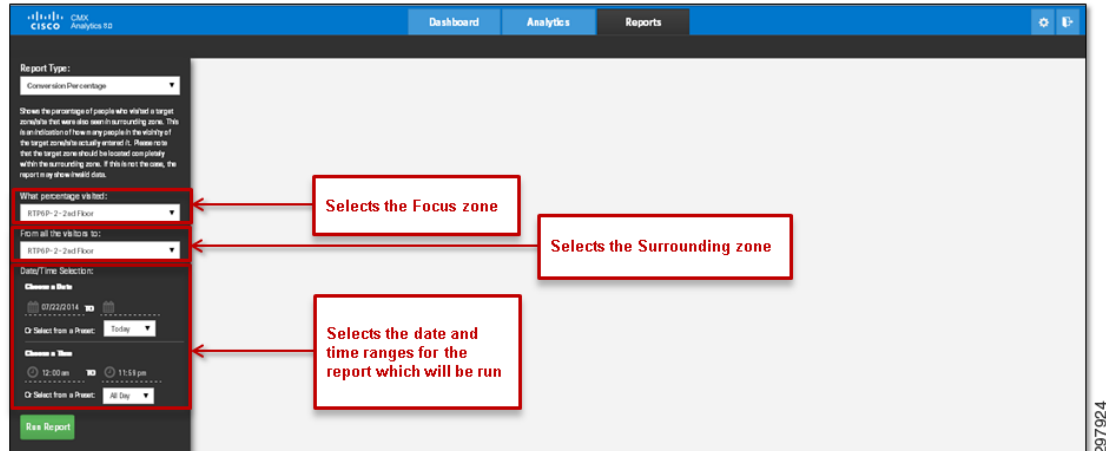
- Conversion Percentage
- Daily Visitors and Dwell Time
- Detected vs. Connected Devices
- Hourly Visitors and Dwell Time
- Movement between Zones
- Repeat Visitors

The following sections briefly explain each report and discuss the configuration options.

### Conversion Percentage Report

The Conversion Percentage Report estimates the percentage of people who were in the vicinity of a zone before entering that zone. An example of the setup screen for running this report is shown in [Figure 26-42](#).

Figure 26-42 Setup Screen for the Conversion Percentage Report

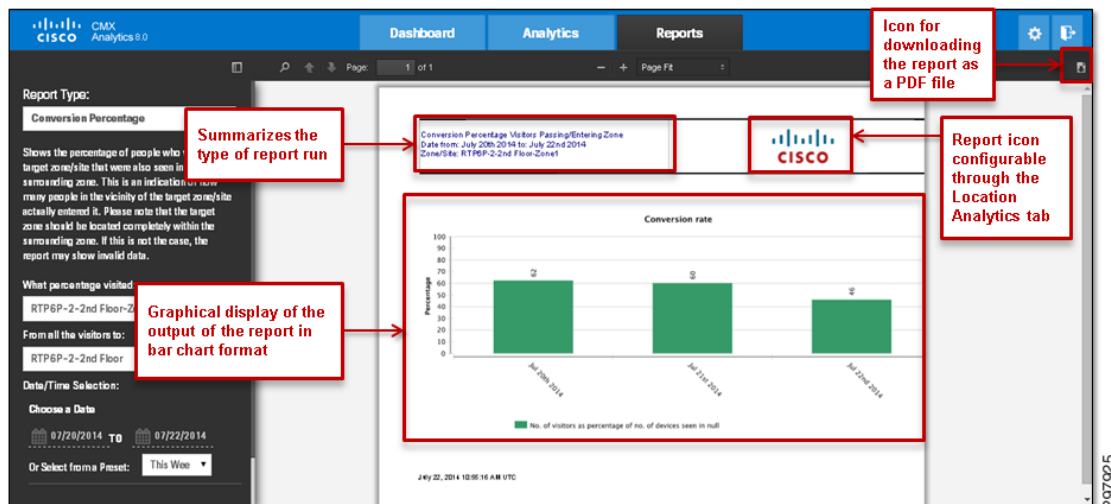


For this report to be effective, a minimum of two zones must be configured within CMX Analytics for the particular site. One zone serves as the focus zone, which is the zone you are interested in where visitors enter, and the second is the catchment or surrounding zone, which contains potential visitors who may or may not enter the focus zone.

The beginning and ending dates can be selected by clicking one of the preset dates: today, yesterday, this week, last week, this month, or last month. Alternatively, the administrator can select the **calendar icons** under **Choose a Date**. The beginning and ending times can be selected by clicking one of the preset times: All Day, Business Hours, Morning, Lunch Time, Afternoon, or Evening. Alternatively, the administrator can select the **clock icons** under **Choose a Time**. By default, the date is set for the current date and the time is set for All Day.

Figure 26-43 provides an example of the output of the Conversion Percentage report, as displayed within the web browser.

Figure 26-43 Example of the Conversion Percentage Report Output Displayed in the Web Browser



The output of each type of report follows a similar format. By default, the output of reports appears as one or more pages within the central panel of the web browser. The top part of the report summarizes the type of report run, as well as the settings which were selected for the report. To the right of that

summary is an icon intended to display a company logo, if the reports are to be viewed online, or exported in .pdf format. The configuration of the icon was discussed in [Configuring CMX Location Analytics](#).

The data within the body of the report is displayed in graphical format. For this particular example, the output shows the percentage of visitors to the entire site, who visited a single zone, in bar chart format for this week. Since only three days have passed in the current week, only three columns are present. Note that the Conversion Percentage report only provides percentages, not absolute numbers. An example of where this type of information may be useful is in a retail store which assumes approximately the same number of customers from week to week. By running this report at the end of each week, the store manager may be able to see if the percentage of customers visiting a particular department—configured as a separate zone within CMX Analytics—is increasing or decreasing over time. This information may be just one additional data point in assessing the effectiveness of marketing campaigns designed to increase foot traffic to a particular department within the store.

Finally, the icon in the upper right corner of the page shown in [Figure 26-43](#) can be used to download reports to the administrator's computer in Adobe Acrobat (.pdf) format. In this format, the output looks the same as displayed within the web browser. These can then be emailed to interested parties for viewing.

## Daily Visitors and Dwell Time Report

The Daily Visitors and Dwell Time Report compares the number of devices in several dwell time categories, for the same time window, over a number of days within the focus zone. An example of the setup screen for running this report is shown in [Figure 26-44](#).

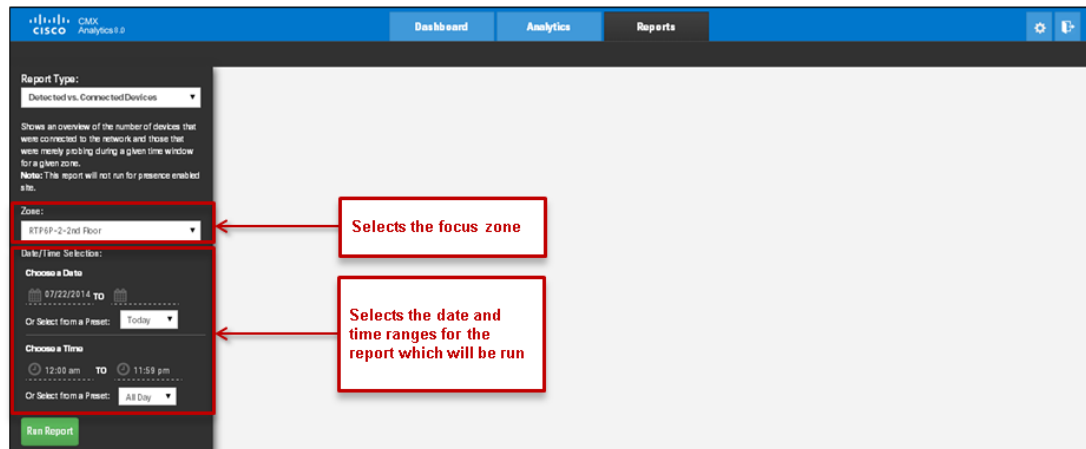
**Figure 26-44** Setup Screen for the Daily Visitors and Dwell Time Report

A single zone, the focus zone, is selected for the report. The beginning and ending dates can be selected by clicking one of the preset dates: today, yesterday, this week, last week, this month, or last month. Alternatively, the administrator can select the **calendar icons** under **Choose a Date**. By default, the date is set for the current date and the time is set for All Day.

## Detected versus Connected Devices Report

The Detected vs. Connected Devices Report shows an overview of the number of devices that were connected to the network (associated with an AP) and the devices that were merely probing during a given time period for a particular zone. An example of the setup screen for running this report is shown in [Figure 26-45](#).

**Figure 26-45** Setup Screen for the Detected vs. Connected Devices Report

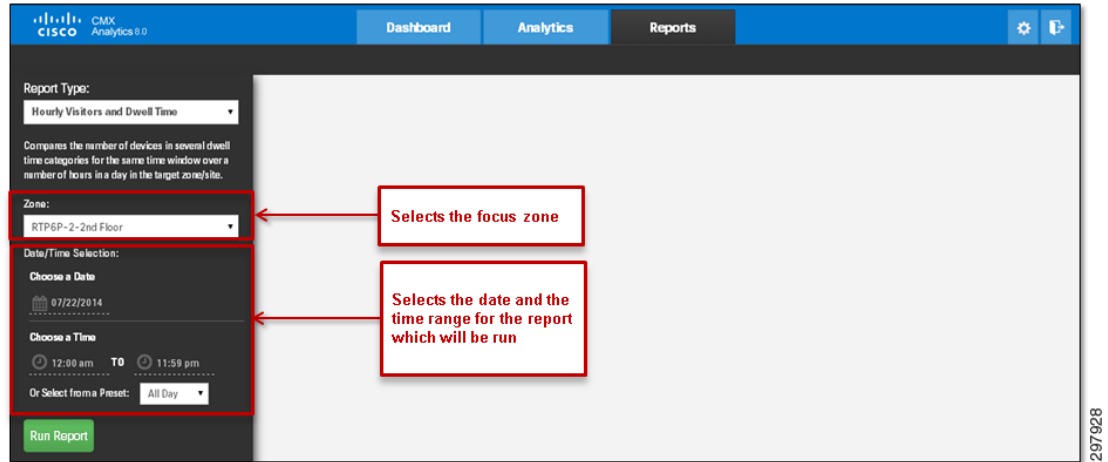


As with the previous report, a single zone, the focus zone, is selected. The beginning and ending dates can be selected by clicking one of the preset dates: today, yesterday, this week, last week, this month, or last month. Alternatively, the administrator can select the **calendar icons** under **Choose a Date**. The beginning and ending times can be selected by clicking one of the preset times: All Day, Business Hours, Morning, Lunch Time, Afternoon, or Evening. By default, the date is set for the current date.

## Hourly Visitors and Dwell Time Report

The Hourly Visitors and Dwell Times Report compares the number of devices in several dwell time categories, for the same time window, over a number of hours in the focus zone. To limit the clutter of results, time windows are only selectable for single days. An example of the setup screen for running this report is shown in [Figure 26-46](#).

**Figure 26-46** Setup Screen for the Hourly Visitors and Dwell Time Report



A single zone, the focus zone, is selected for the report. The date can be selected by clicking the calendar icon under Choose a Date. The beginning and ending times can be selected by clicking one of the preset times: All Day, Business Hours, Morning, Lunch Time, Afternoon, or Evening. Alternatively, the administrator can select the **clock icons** under **Choose a Time**. By default, the date is set for the current date and the time is set for All Day.

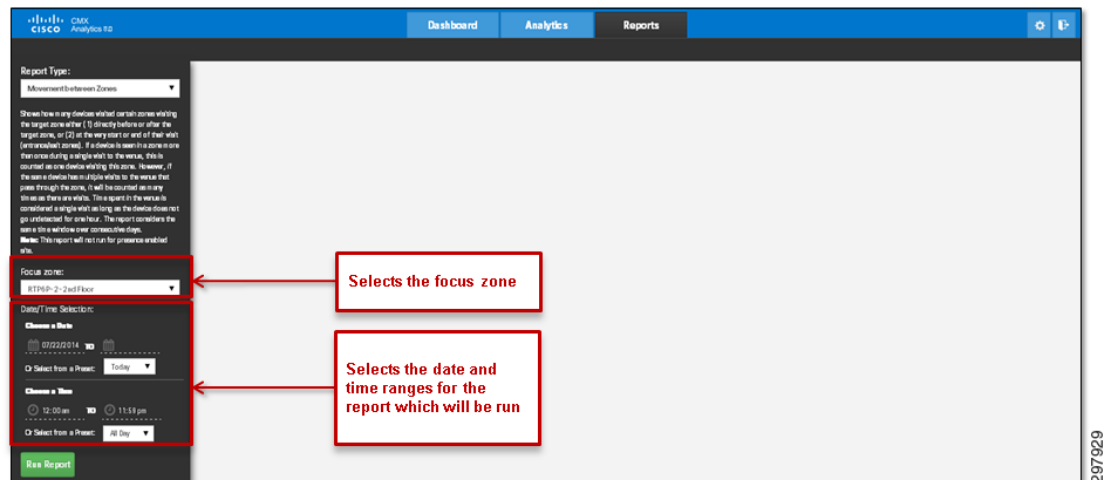
## Movement between Zones Report

The Movement between Zones Report provides a breakdown of all zones at specific points as devices pass to and from the focus zone. This report gives the following information:

- The zone where the device was first detected.
- The immediate zone before and after where the device was last detected.

An example of the setup screen for running this report is shown in [Figure 26-47](#).

**Figure 26-47** Setup Screen for the Movement between Zones Report

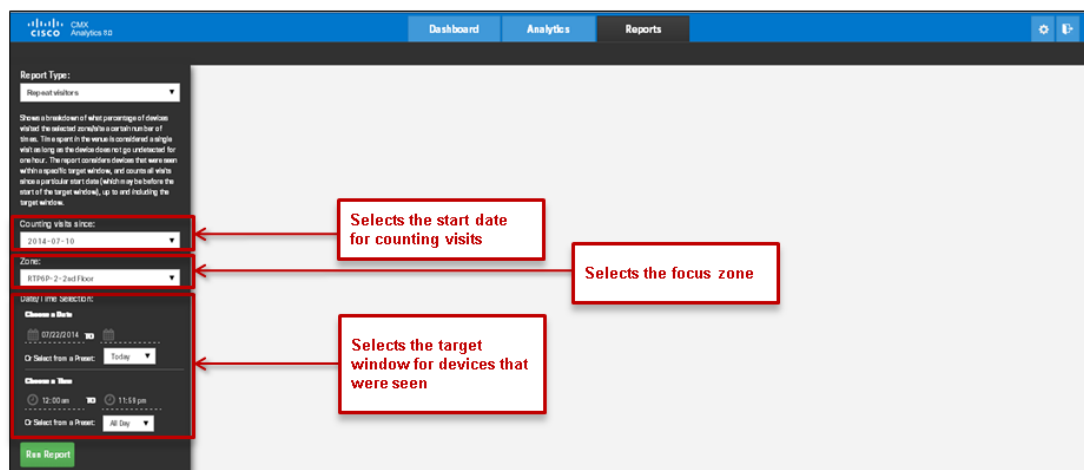


A single zone, the target zone, is selected for the report. The beginning and ending dates can be selected by clicking one of the preset dates: today, yesterday, this week, last week, this month, or last month. Alternatively, the administrator can select the **calendar icons** under **Choose a Date**. The beginning and ending times can be selected by clicking one of the preset times: All Day, Business Hours, Morning, Lunch Time, Afternoon, or Evening. Alternatively, the administrator can select the **clock icons** under **Choose a Time**. By default, the date is set for the current date and the time is set for All Day.

## Repeat Visitors Report

The Repeat Visitors Report shows a breakdown of what percentage of devices visited a selected zone a certain number of times. Time spent in the venue is considered a single visit as long as the device does not go undetected for one hour. An example of the setup screen for running this report is shown in Figure 26-48.

**Figure 26-48** Setup Screen for the Repeat Visitors Report



The report considers devices that were seen within a specific target window and counts all visits since a particular start date (which may be before the start of the target window), up to and including the target window.

The CMX administrator must configure three parameters. First, they must select the focus zone for the report. Next, the CMX administrator must select a target window. The beginning and ending dates can be selected by clicking one of the preset dates: today, yesterday, this week, last week, this month, or last month. Alternatively, the administrator can select the **calendar icons** under **Choose a Date**. The beginning and ending times can be selected by clicking one of the preset times: All Day, Business Hours, Morning, Lunch Time, Afternoon, or Evening. Alternatively, the administrator can select the **clock icons** under **Choose a Time**. By default, the date is set for the current date and the time is set for All Day. Only devices seen within the target window are counted within the output of the report. Finally, the CMX administrator must select a start date. This is when the counting of repeat visits of the devices seen within the target window begins. The start date may be before the start of the target window.



## Configuring CMX Visitor Connect

---

September 4, 2014

This chapter highlights the configuration options available for CMX Visitor Connect, including creating splash pages and social connectors and applying them to floors, as well as configuring a Facebook social connector. Finally the configuration of Role Based Access Control (RBAC) for the overall CMX Connect & Engage service is discussed.



**Note**

---

The configuration of CMX Facebook Wi-Fi, which is a separate feature of the CMX Connect & Engage service, is not covered in the current version of the CMX CVD.

---

## Configuring CMX Visitor Connect with Splash Pages and Social Connectors

To configure the CMX Visitor Connect service with a splash page and a Facebook social connector:

- 
- Step 1** From the MSUI dashboard, click the **Connect & Engage icon** to be taken to the CMX Connect & Engage UI. Alternatively you can directly go to the CMX Connect & Engage UI by going to the URL:  
`https://<MSE_Name_or_IP_Address>/dashboard/`  
The default username is admin and the default password is admin, which should be changed. [Configuring RBAC on CMX Connect & Engage](#) discusses the configuration of additional groups and users for the CMX Connect & Engage service on the MSE.
- Step 2** Click **Visitor Connect > Splash Templates** to display the Splash Templates Configuration screen, as shown in [Figure 27-1](#).

Figure 27-1 Splash Template Configuration Screen

The screenshot displays the 'Splash Template Configuration' interface. At the top, there's a blue header with the Cisco logo and 'CMX Connect & Engage'. A sidebar on the left contains navigation items: Summary, Visitor Connect (with 'Splash Templates' highlighted), Template Fields, Social Connectors, Facebook Wi-Fi, Visitor Policy, Maps, Mobile App, Accounts, and Settings. The main content area is titled 'Splash Template Configuration' and features a 'Splash Templates' section. This section includes a four-step process: STEP 1 (Optional) Create Template Fields, STEP 2 (Optional) Create Social Connectors, STEP 3 Create Splash Templates, and STEP 4 Assign the Splash Template. Below the steps, there's explanatory text about how MSE selects templates based on location granularity and a hierarchy of Zone -> Floor -> Venue -> Campus. At the bottom, there's a 'Splash Templates' table with columns for Name, Default Template, Facebook, LinkedIn, and Google+. A red arrow points to the '+ Create' button, and a callout box says 'Click Create to create a new splash template'. The table currently shows 'No data available in table' and 'Showing 0 to 0 of 0 entries'. The footer includes 'Powered by Cisco. All Rights Reserved.' and a vertical ID '297981'.

Existing templates are listed in the Splash Templates section of the screen. These can be edited by highlighting the template and clicking the **Edit** button. Click the **Create** button to create a new splash template, as shown in Figure 27-1.

**Step 3** Figure 27-2 shows the available fields within the Add/Edit Splash Template screen.



Figure 27-2 Add/Edit Splash Template Screen

These fields include:

- **Template Name**—Multiple templates can be supported by CMX Visitor Connect. Enter a descriptive name for the template to distinguish it from other templates.
- **Background**—From the drop down menu adjacent to the Background field, select one of the pre-configured background images for the template. Alternatively, you can upload your own image to be used as the background for the splash page.
- **Header**—You can create a custom header which appears at the top of the splash page. For example, you can welcome guests to the venue via a simple text message. The appearance of the text can be customized by adjusting the font, font size, color, etc. Clicking the **Source** tab to configure the header directly in HTML markup if desired.
- **Form Fields**—The operator of the venue can optionally collect information from visitors in exchange for the use of the free Wi-Fi service. To use a form field within a template, it must first be created by clicking the **Create Template Fields** button. This displays the Create Template Field popup window, as shown in [Figure 27-3](#).

**Figure 27-3** Create Template Field Popup Window

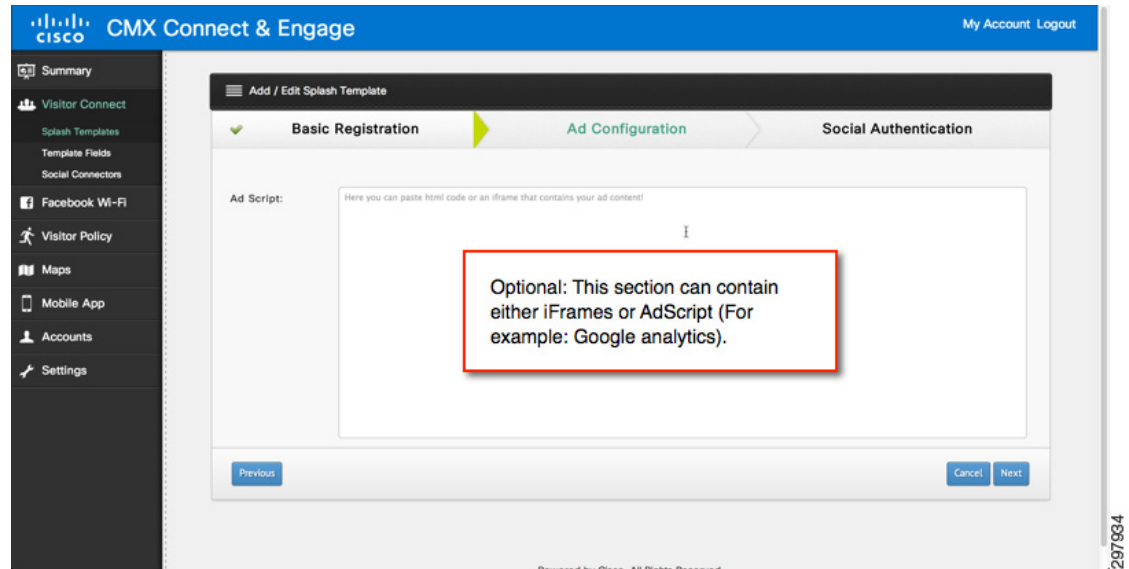
Since multiple template fields can be added to a template, select a unique but obvious name for the template field which asks for the information to be collected. The venue operator can select **Text** if the guest is asked to fill in a simple text field. If the guest is asked to fill in a field with comma-separated values, the venue operator can select **List**. Once the form field is created by selecting the **Create Field** button, it appears in the drop-down menu of existing template fields under **Pick from existing Template Fields**, as shown in [Figure 27-2](#).



**Note** Alternatively, the venue operator can follow the steps shown in the configuration wizard and create Template Fields prior to configuring Splash Templates. This is done by selecting **Visitor Connect > Template Fields** from the Splash Template Configuration screen, as shown in [Figure 27-1](#).

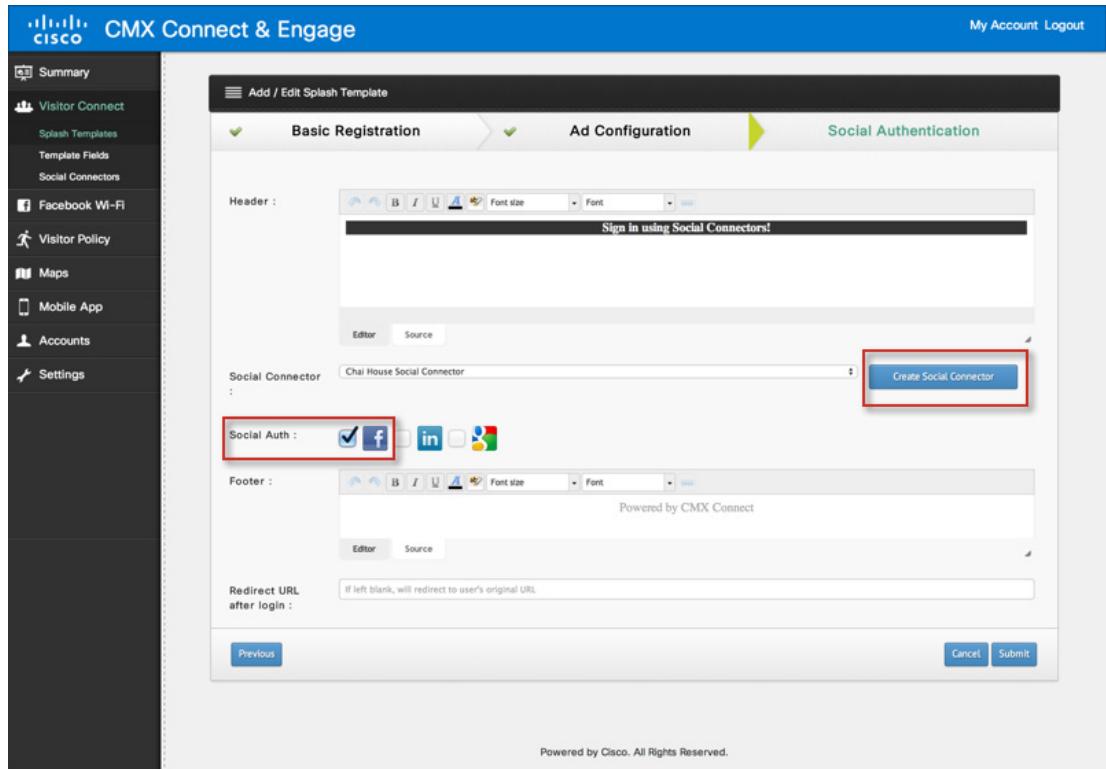
- **Terms & Conditions**—This field can be used to generate a text field showing the guest, terms, and conditions for use of the Wi-Fi service within the venue. As with the Header field, the appearance of the text can be customized by adjusting the font, font size, color, etc. Click the **Source** tab to configure the terms and conditions directly in HTML markup if desired.
- **Footer**—You can also create a customer footer which appears at the bottom of the splash page. Again, the appearance of the text can be customized by adjusting the font, font size, color, etc. Click the **Source** tab to configure the footer directly in HTML markup if desired.

Once you have entered the details for the splash page, click the **Next** button at the bottom of the page. This takes you to the Add Configuration screen, as shown in [Figure 27-4](#).

**Figure 27-4** Splash Page Ad Configuration

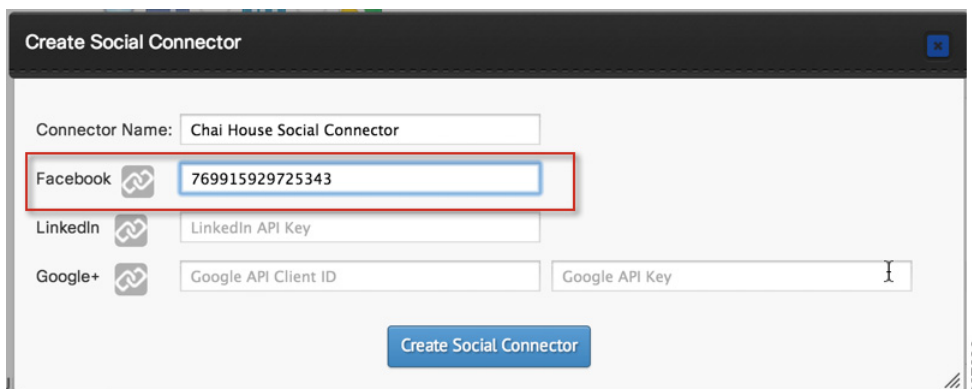
- Step 4** Optionally, enter an Ad Script in the field provided. The Ad Script is an optional step which can be used to position marketing content (which may be located on a separate server) to the visitor, if desired. Clicking the **Next** button takes you to the social authentication step.
- Step 5** To configure Social Authentication, in [Figure 27-5](#) click the **Create Social Connector** button to display a dialog box (popup window).

Figure 27-5 Splash Page Social Authentication



- Step 6** An example of creating a Facebook social connector is shown in [Figure 27-6](#). In the Create Social Connector popup window, give a name to the social connector and enter the Facebook App id. The process of creating a Facebook App for social authentication is explained in [Configuring Facebook App for Visitor Connect](#).

Figure 27-6 Create Social Connector



Click the **Create Social Connector** button to go back to the Social Authentication page.

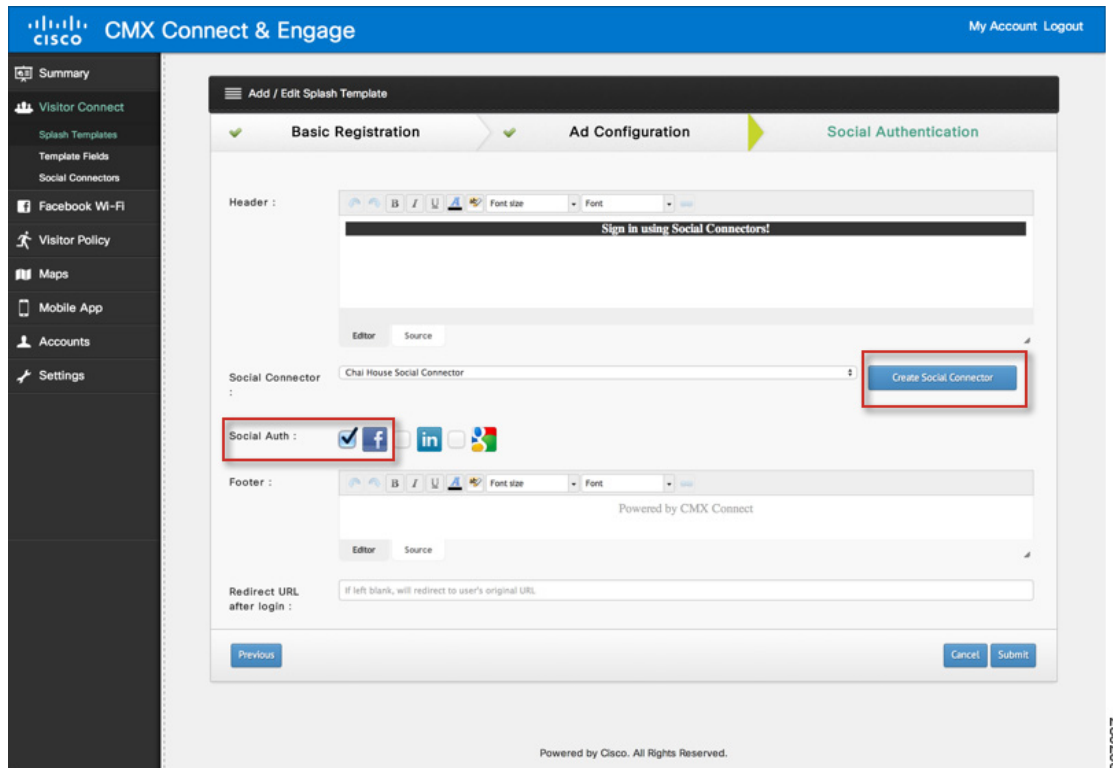


**Note** Alternatively, you can follow the steps shown in the configuration wizard and create social connectors prior to configuring splash templates. This is done by selecting **Visitor Connect > Social Connectors** in the Splash Template Configuration screen shown in [Figure 27-1](#).

You can also create a customer header and footer which appear at the top and bottom of the page when the visitor is redirected to authenticate to social media. The appearance of the text can be customized by adjusting the font, font size, color, etc. Clicking the **Source** tab to configure the header or footer directly in HTML markup if desired.

- Step 7** To enable Facebook social authentication for the splash page, select the checkbox next to Facebook which is adjacent to the Social Auth: field, as shown in [Figure 27-7](#).

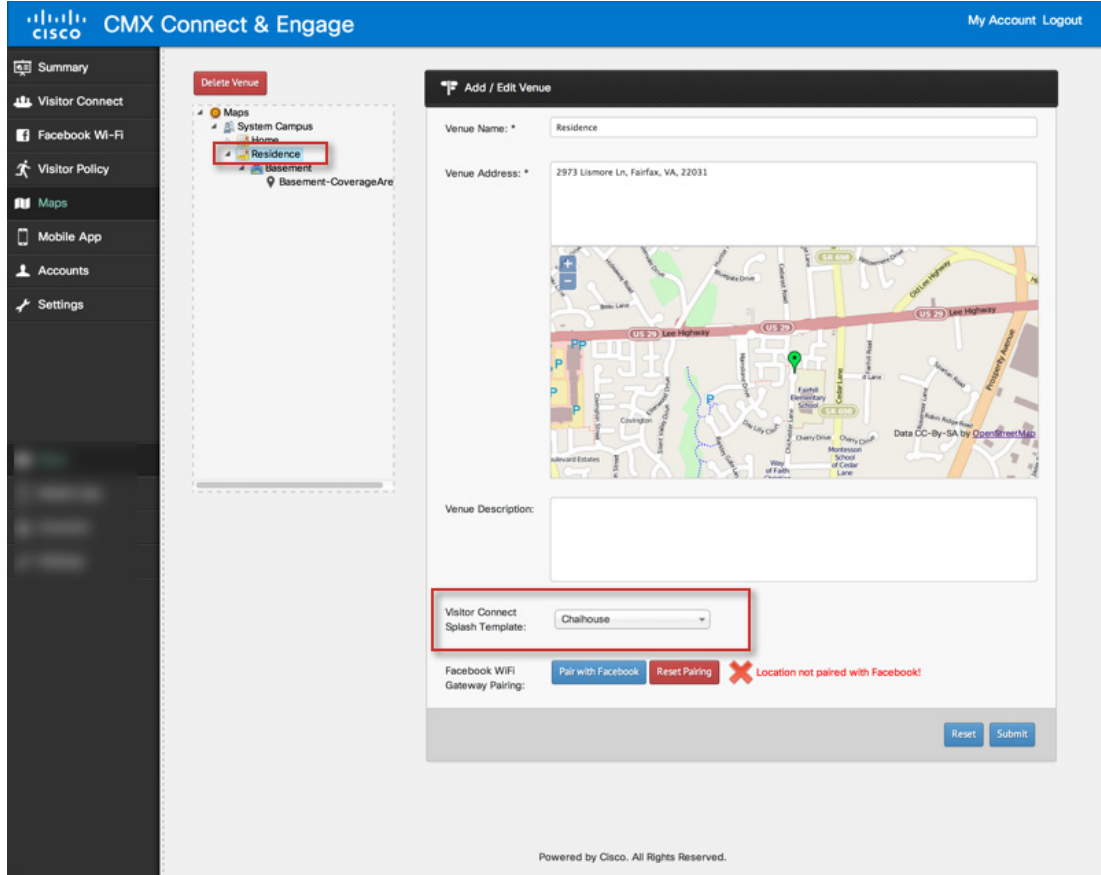
**Figure 27-7** Enable Social Authentication



Click the **Submit** button at the bottom of the screen to save the changes.

- Step 8** Once the splash template has been saved, it can be applied. Splash templates can be applied across a campus, within a venue (which corresponds to a building), within a floor, or within a zone. To do this, select **Maps**, located on the left panel of the CMX Connect & Engage service. This displays the Maps screen, as shown in [Figure 27-8](#).

Figure 27-8 Apply Splash Templates to Maps




Expand the maps tree until the particular campus, venue, or floor is visible and select it. In the drop-down menu next to Visitor Connect Splash Template, select the appropriate splash template. Click the **Submit** button to apply the template.

## Configuring Facebook App for Visitor Connect

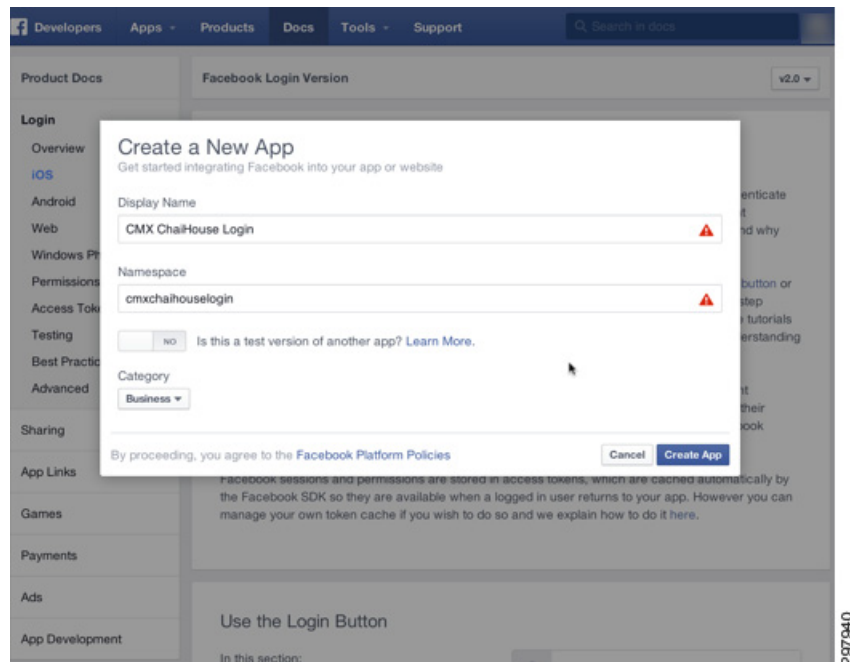
In this section we create an example Facebook application that can be used as a social authenticator with a splash page. While it is possible to use a Google+ app login and/or LinkedIn app, they are not covered as part of this version of the CMX guide. To get started with a Facebook App, log in to the Facebook developer portal at the following URL with your Facebook ID:

<http://developer.facebook.com>

**Step 1** Click **Apps > Create a New App** to get started, as shown in [Figure 27-9](#).

**Figure 27-9** Create a Facebook App

**Step 2** The Create a New App popup window is displayed, as shown in [Figure 27-10](#).

**Figure 27-10** Facebook App Details

Give a descriptive name to your new App. Also create a namespace for the App by making sure there are no spaces in the name.

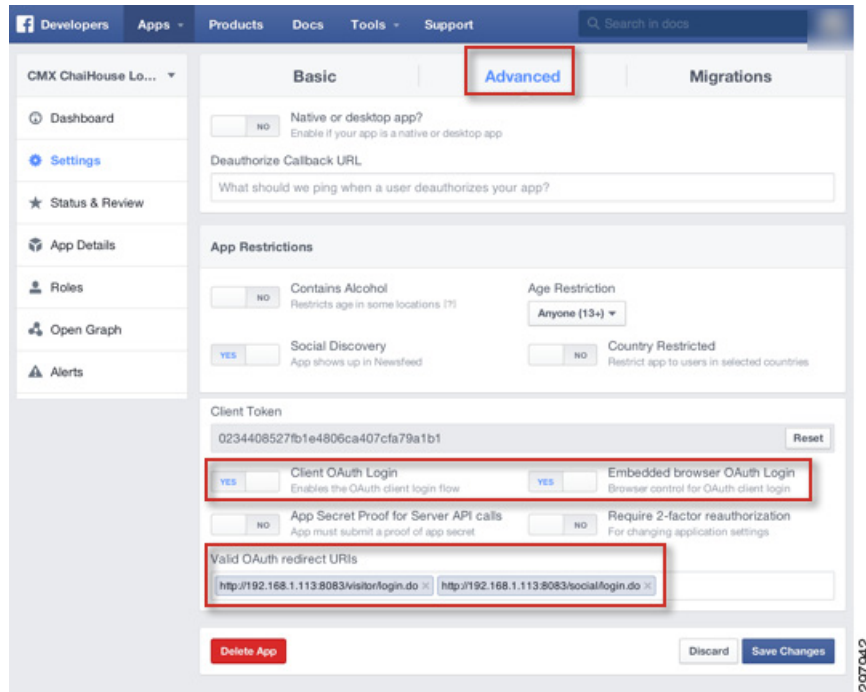
**Step 3** Once the App is created, make a note of the App ID. This is used as the Facebook App ID when configuring the social connector while creating a splash page template within CMX Visitor Connect. An example of the App ID is shown in [Figure 27-11](#).

**Figure 27-11** Getting the Facebook App ID

Figure 27-6 in [Configuring CMX Visitor Connect with Splash Pages and Social Connectors](#) showed where the Facebook App ID is applied within CMX Visitor Connect.

**Step 4** Click the **Advanced** tab for the App, as shown in [Figure 27-12](#).

**Figure 27-12 Advanced App Configuration**



Ensure that both Client OAuth Login and Embedded browser OAuth Login are enabled and in the “Yes” state. Under Valid OAuth Redirect URIs, enter the following URLs in the fields respectively:

- `http://< MSE_IP_Address>:8083/visitor/login.do`
- `http://< MSE_IP_Address>:8083/social/login.do`
- MSE\_IP\_Address corresponds to the IP address of the MSE.

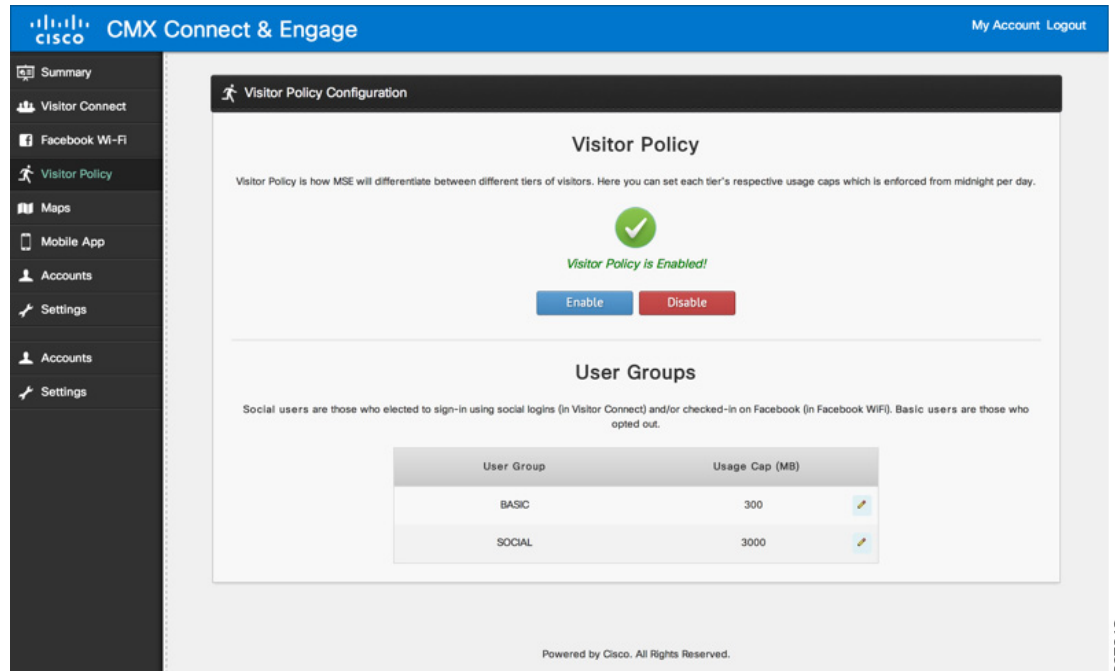
This completes setting up the Facebook App to work with CMX Visitor Connect.

## Visitor Policy

The Visitor Policy section of CMX Connect & Engage allows an administrator to set different usage quotas to users based on the User Group to which the visitor belongs. The User Group, in turn, corresponds to the login type the visitor used when connecting via CMX Visitor Connect. By default, all visitors who do not login with a social connector via the CMX Visitor Connect splash page are put in the Basic group. The Basic group is allocated a usage quota of 300 Mbytes of data for each 24 hour period. Visitors who choose to login via a social connector are put in to Social group and allocated a usage quota of 3000 Mbytes of data for each 24 hour period. This can be used to encourage visitors to use social connectors with splash pages. The usage quotas may be easily edited by clicking the **Visitor Policy** tab on the CMX Connect & Engage dashboard, as shown in [Figure 27-13](#).



Figure 27-13 Visitor Policy Configuration

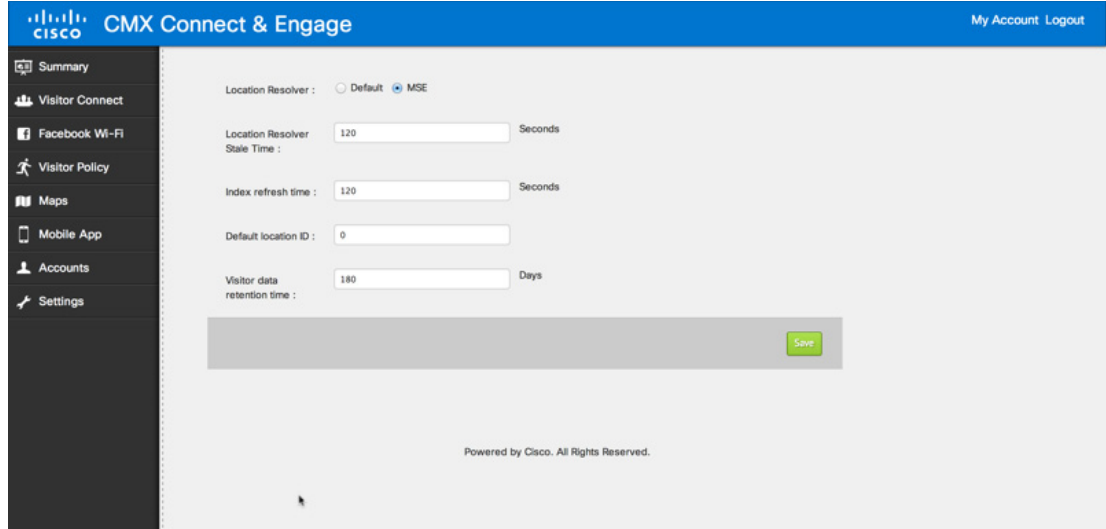


The administrator can also choose to enable or disable the use of usage quotas entirely within the Visitor Policy screen.

## Server Settings

Server Settings allow you to configure retention time for visitor data. Server Settings can be accessed by clicking **Settings** > **Server Settings** from the CMX Connect & Engage page. The default retention time for visitor data is 180 days and can be configured globally according to the administrator's requirements. [Figure 27-14](#) shows an example of the Server Settings screen within CMX Connect & Engage.

Figure 27-14 Connect &amp; Engage Server Settings



## Configuring RBAC on CMX Connect & Engage

The CMX Connect & Engage service has its own role-based access control (RBAC) separate from the MSE itself. For role-based access control on the MSE, see [Configuring Role-Based Access Control \(RBAC\) on the MSE](#) in [Chapter 25, “Configuring the Mobility Services Engine for CMX.”](#)

CMX Connect & Engage provides very granular role-based access control (RBAC). Each role can be configured for each of the following 15 operations:

- Accounts—Allows members of the role to create, edit, and delete Accounts. Accounts are associated with different Campaigns and Banners.
- Banner Approver—Allows members of the role to approve Banners.
- Banners—Allows members of the role to create and edit Banners.
- Campaigns—Allows members of the role to create and edit Campaigns.
- Campaign Approver—Allows members of the role to approve Campaigns.
- CMX Mobile—Allows members of the role to access functions found under the Mobile App topic.
- Domain Setup—Allows members of the role to create, edit, and delete domains, which are found under the Settings topic.
- Floor Navigation—Allows members of the role to access Floor Navigation functions found under the Mobile App topic.
- Menu
- Point of Interest—Allows members of the role to create, edit, and delete points of interest.
- Reports
- Roles—Allows members of the role to create, edit, and delete Roles.
- Server Settings—Allows members of the role to access Server Settings functions found under the Settings topic.

- Users—Allows members of the role to create, edit, and delete Users. Users are associated to a particular role for role-based access control.
- Visitor Connect—Allows members of the role to create splash templates and configure social media connectors.

To configure RBAC, the CMX administrator must first login to the CMX Connect & Engage service running on the MSE via the graphical user interface (GUI).


The following provides an example of the URL to access the CMX Connect & Engage Summary page.

`https://<MSE_IP_Address>/dashboard/`

MSE\_IP\_Address is the IP address of the MSE server which is running the CMX Connect & Engage service.

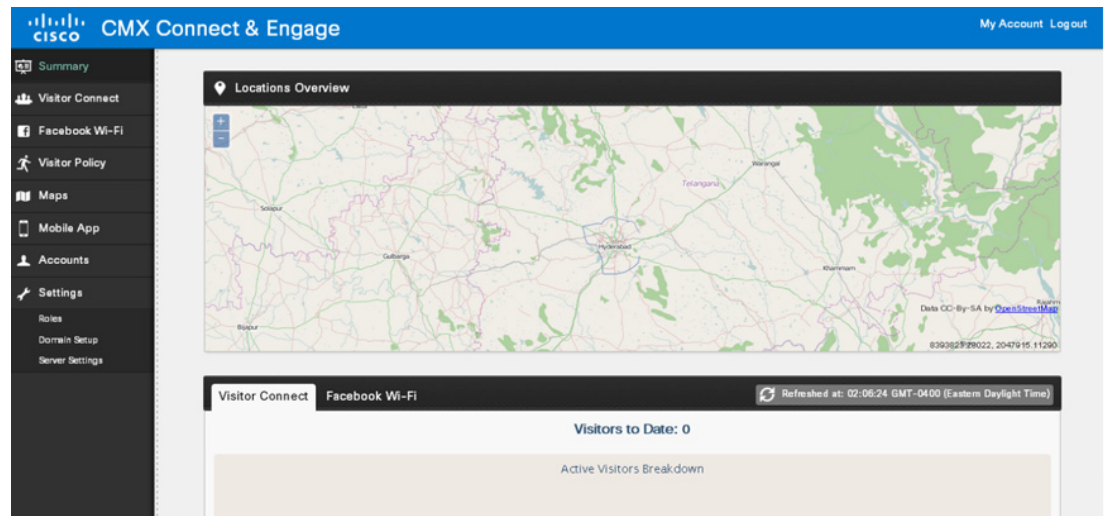
Figure 27-15 is an example of the login screen which should be displayed.

**Figure 27-15** CMX Connect & Engage Login Page



Upon logging in, the CMX administrator is automatically taken to the CMX Connect & Engage Summary page, as shown in Figure 27-16.

**Figure 27-16** Example of CMX Connect & Engage Summary Page



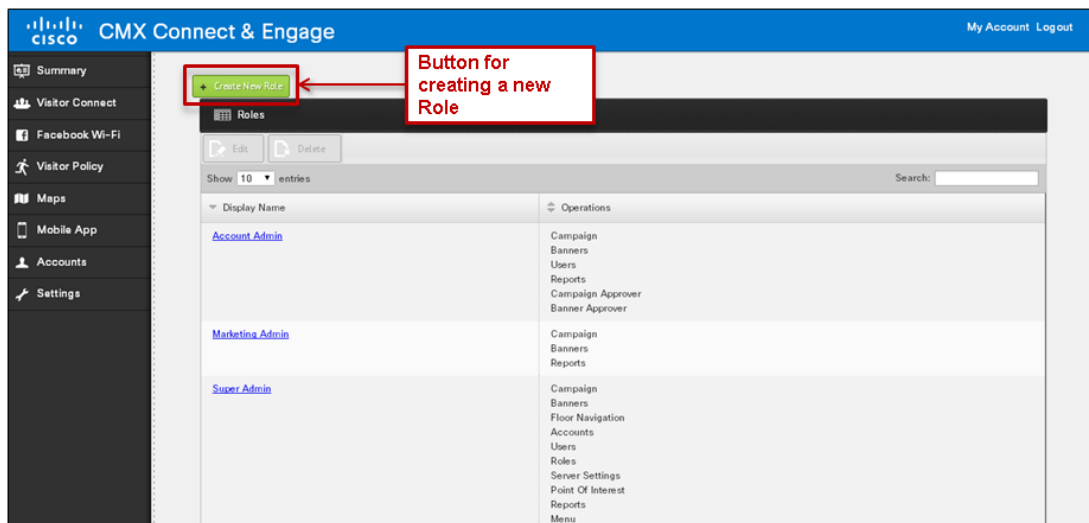
The panel on the left side of the page has eight main topics for configuration of the CMX Connect & Engage service:

- Summary

- Visitor Connect
- Facebook Wi-Fi
- Visitor Policy
- Maps
- Mobile App
- Accounts
- Settings

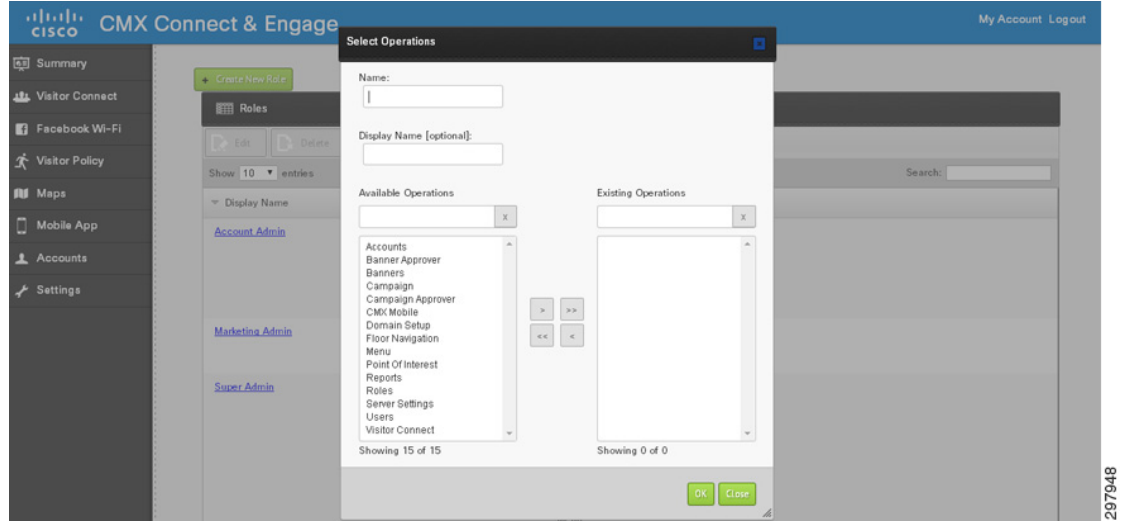
The CMX administrator must first configure one or more Roles by clicking the **Settings** topic on the panel on the left side of the page and selecting **Roles** from the drop-down menu which is displayed. This displays the Roles page, as shown in [Figure 27-17](#).

**Figure 27-17** Example of the Roles Page



When the CMX administrator clicks the **Create New Role** button, a popup page is displayed, as shown in [Figure 27-18](#).

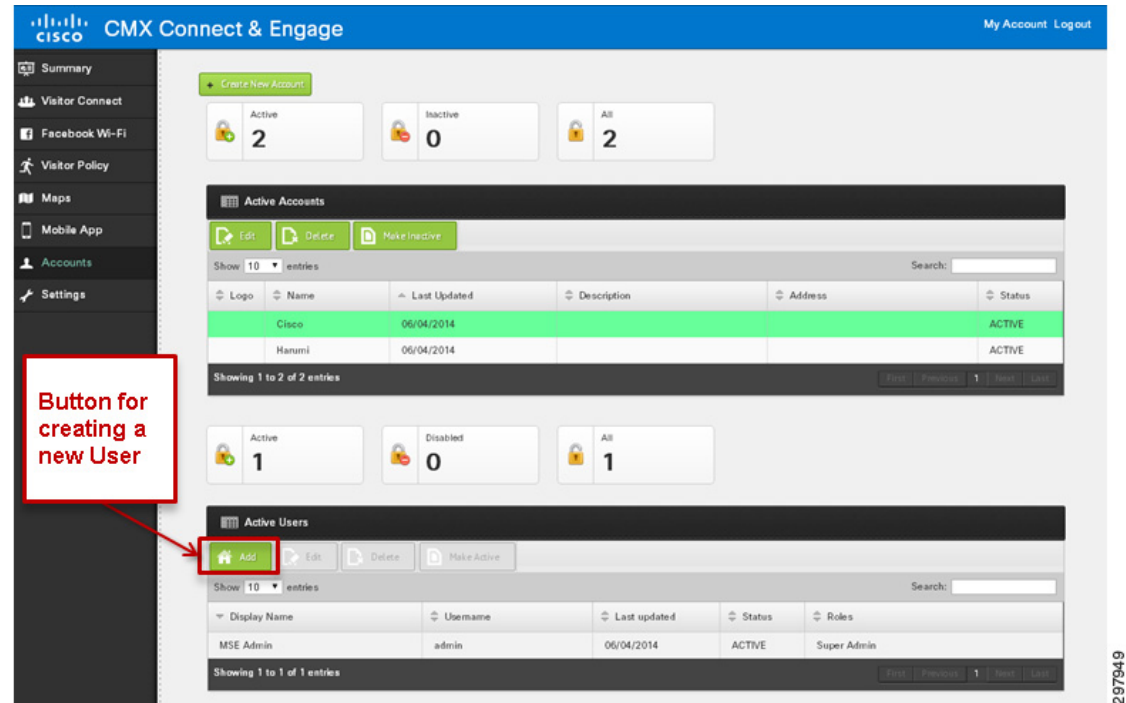
Figure 27-18 Creating a New Role



The popup window has fields for the Name of the role and an optional Display Name. The CMX administrator can then select from the available 15 operations discussed above to assign to the new role. Clicking the **OK** button adds the new role with selected operations. Clicking **Close** cancels the addition of the new role.

Once the CMX administrator has added the new role, they can add individual users to the role by clicking the **Accounts** topic located in the panel on the left side of the main CMX Connect & Engage page. This displays the page where accounts and users are created, as shown in Figure 27-19.

Figure 27-19 Example of the Page to Create Accounts and Users



When the CMX administrator clicks the **Add** button (under the Active Users section toward the bottom of the page), the Add/Edit Users page is displayed, as shown in [Figure 27-20](#).

**Figure 27-20** Example of the Add/User Page

The Add/Edit User page has mandatory fields for the Display Name, Username, Password, Confirm Password, and E-mail address of the new user. The CMX administrator can also select the role to which the new user will belong from the drop-down menu in the Roles field. By default, new users are created in a status of DISABLED. To enable the new user, the CMX administrator must select the **Make Active** button

Clicking the **Submit** button adds the new user with desired role. Clicking **Cancel** cancels the addition of the new user.



## **PART 6**

### **CMX Appendices**







# CMX Software Versions

Revised: September 4, 2014

The following tables highlight the hardware and software components used for validation testing in this design guide. While other software versions may provide the required CMX functionality, the Validated Software Version column indicates the software version used during the validation. For specific feature support, consult the Cisco Feature Navigator: <http://www.cisco.com/go/fn>.

**Table A-1** *CMX Infrastructure Components*

Component	Role	Validated Software Version
Cisco Mobility Services Engine	Mobility Services Engine	MSE 8.0.100.0
Cisco 5508 Wireless Controller	Wireless LAN Controller	WLC 8.0.100.0
Cisco Flex 7500 Wireless Controller	Wireless LAN Controller	WLC 8.0.100.0
Cisco AP 3702i	Access Point	N/A
Cisco AP 3602i	Access Point	N/A
Cisco AP 2602i	Access Point	N/A
Cisco Wireless Security Module	Wireless Security Module	N/A

**Table A-2** *Mobile Devices*

Device	Validated Software Version
Apple Macbook Pro	OSX 10.7.5
Samsung Galaxy S3	Android OS 4.3
Apple iPhone 5S	iOS 6.1.4
Apple iPad Mini	iOS 7.0.6
Apple iPad Air	iOS 7.1.1
Samsung Galaxy S2	Android OS 3.2
HTC EVO (cell)	Android OS 2.3.5
Lenovo PC	Windows 7 with 802.11g adapter





## CMX System Release Notes

---

Revised: September 4, 2014

### MSE 8.0 Role-Based Access Control

As of MSE version 8.0 it is not recommended to utilize Read Access groups on the MSE. CMX Presence Analytics currently does not participate in Role-Based Access Control (RBAC) on the MSE. Hence any userid which is part of a Read Access group will also be able to add/modify/delete CMX Presence Analytics configuration. Note that the use of only Write Access and Full Access groups means that any non-IT personnel who require access to the CMX Analytics dashboard, analytics tab, or reports may also have access to add/modify/delete CMX Analytics (location and presence) configuration. One way to mitigate some of this risk is for IT personnel to download CMX Analytics Reports and email them to non-IT personnel at regular intervals. Alternatively, the list of non-IT personnel—such as store operations managers, marketing executives, etc.—who have direct access to the MSE for CMX Analytics should be kept tightly controlled.





## 802.11 Data Rates

Revised: September 4, 2014

This appendix lists data rates for 802.11an/ac rates in 5 GHz and 802.11bgn in 2.4GHz.

### IEEE 802.11a/n/ac

The channel identifiers, channel center frequencies, and regulatory domains of each IEEE 802.11a 20-MHz-wide channel are listed in [Table C-1](#).

**Table C-1** Channels for IEEE 802.11a

Channel Identifier	Frequency in MHz	Regulatory Domains			
		Americas (-A)	Japan (-J)	Singapore (-S)	Taiwan (-T)
34	5170	-	X	-	-
36	5180	X	-	X	-
38	5190	-	X	-	-
40	5200	X	-	X	-
42	5210	-	X	-	-
44	5220	X	-	X	-
46	5230	-	X	-	-
48	5240	X	-	X	-
52	5260	X	-	-	X
56	5280	X	-	-	X
60	5300	X	-	-	X
64	5320	X	-	-	X
149	5745	-	-	-	-
153	5765	-	-	-	-
157	5785	-	-	-	-
161	5805	-	-	-	-

**Note**

All channel sets are restricted to indoor usage except the Americas (-A), which allows for indoor and outdoor use on channels 52 through 64 in the United States.

## IEEE 802.11b/g/n

The channel identifiers, channel center frequencies, and regulatory domains of each IEEE 802.11b 22-MHz-wide channel are listed in [Table C-2](#).

**Table C-2** Channels for IEEE 802.11b

Channel Identifier	Frequency in MHz	Regulatory Domains				
		Americas (-A)	EMEA (-E)	Israel (-I)	China (-C)	Japan (-J)
1	2412	X	X	-	X	X
2	2417	X	X	-	X	X
3	2422	X	X	X	X	X
4	2427	X	X	X	X	X
5	2432	X	X	X	X	X
6	2437	X	X	X	X	X
7	2442	X	X	X	X	X
8	2447	X	X	X	X	X
9	2452	X	X	X	X	X
10	2457	X	X	-	X	X
11	2462	X	X	-	X	X
12	2467	-	X	-	-	X
13	2472	-	X	-	-	X
14	2484	-	-	-	-	X

**Note**

Mexico is included in the Americas regulatory domain; however channels 1 through 8 are for indoor use only while channels 9 through 11 can be used indoors and outdoors. Users are responsible for ensuring that the channel set configuration complies with the regulatory standards of Mexico.

**Note**

France is included in the EMEA regulatory domain; however only channels 10 through 13 can be used in France. Users are responsible for ensuring that the channel set configuration complies with the regulatory standards of France.

# Maximum Power Levels and Antenna Gains

## IEEE 802.11a

An improper combination of power level and antenna gain can result in Effective Isotropic Radiated Power (EIRP) above the amount allowed per regulatory domain. [Table C-3](#) indicates the maximum power levels and antenna gains allowed for each IEEE 802.11a regulatory domain.

**Table C-3** Maximum Power Levels Per Antenna Gain for IEEE 802.11a

Regulatory Domain	Maximum Power Level (mW) with 6-dBi Antenna Gain
Americas (-A) (200 mW to 800 mw, see below)	40
Japan (-J) (10 mW/MHz EIRP maximum)	40
Singapore (-S) (100 mW EIRP maximum)	20
Taiwan (-T) (800 mW EIRP maximum)	40

In addition, Americas (-A) domain regulation provide different maximum power level for channels as shown in [Table C-4](#).

**Table C-4** Maximum Transmit Power IEEE 802.11a

Regulatory Domain	EIRP Maximum
UNII-1 Low Indoor	200 mW
UNII-2 Mid	1 W
UNII Indoor/Outdoor DFS	1 W
UNII-3 Upper	200 W

## IEEE 802.11b

An improper combination of power level and antenna gain can result in EIRP above the amount allowed per regulatory domain. [Table C-5](#) indicates the maximum power levels and antenna gains allowed for each IEEE 802.11b regulatory domain.

**Table C-5** Maximum Power Levels Per Antenna Gain for IEEE 802.11b

Regulatory Domain	Antenna Gain (dBi)	Maximum Power Level (mW)
Americas (-A) (4 W EIRP maximum)	0	100
	2.2	100
	5.2	100
	6	100
	8.5	100
	12	100
	13.5	100
	21	20
EMEA (-E) (100 mW EIRP maximum)	0	100
	2.2	50
	5.2	30
	6	30
	8.5	5
	12	5
	13.5	5
	21	1
Israel (-I) (100 mW EIRP maximum)	0	100
	2.2	50
	5.2	30
	6	30
	8.5	5
	12	5
	13.5	5
	21	1
China (-C) (10 mW EIRP maximum)	0	5
	2.2	5
	5.2	n/a
	6	n/a
	8.5	n/a
	12	n/a
	13.5	n/a
	21	n/a



**Table C-5** *Maximum Power Levels Per Antenna Gain for IEEE 802.11b*

Japan (-J) (10 mW/MHz EIRP maximum)	0	50
	2.2	30
	5.2	30
	6	30
	8.5	n/a
	12	n/a
	13.5	5
	21	n/a





## CMX Use Case Example—Upgrade VoWLAN Ready Network to Location/CMX Ready

---

**Revised: September 4, 2014**

As an example, let us examine a Voice access point layout for the 275 x 159 foot facility. This space is a drywall office and indoor commercial office environment with a path loss exponent of 3.5. These access point locations were selected based on desired signal strength and overlap calculations that were performed by the original designer. In architecting this design, the designer's intention was to provide a solution that closely followed Cisco VoWLAN design best practices.

We opted for a dual band infrastructure, with an 802.11an/ac 5 GHz WLAN that is used by VoWLAN handsets and high-speed WLAN client devices. 802.11bgn 2.4 GHz operations is also supported, but due to the substantially reduced overall capacity on 802.11bgn brought about by the existence of only three non-interfering channels, its use is restricted to legacy data and voice devices. Legacy data devices would include devices that are unable to migrate to 802.11an/ac for reasons such as the client hardware device being no longer offered for sale, battery life concerns, and so on. Candidate legacy devices might include PDAs, older smartphones and tablets, and other devices with embedded wireless onboard that is not easily upgradeable. In the case of our example, we assume that there are still some users of 802.11bgn voice devices present in the environment that have not yet been addressed with 802.11an/ac replacements.



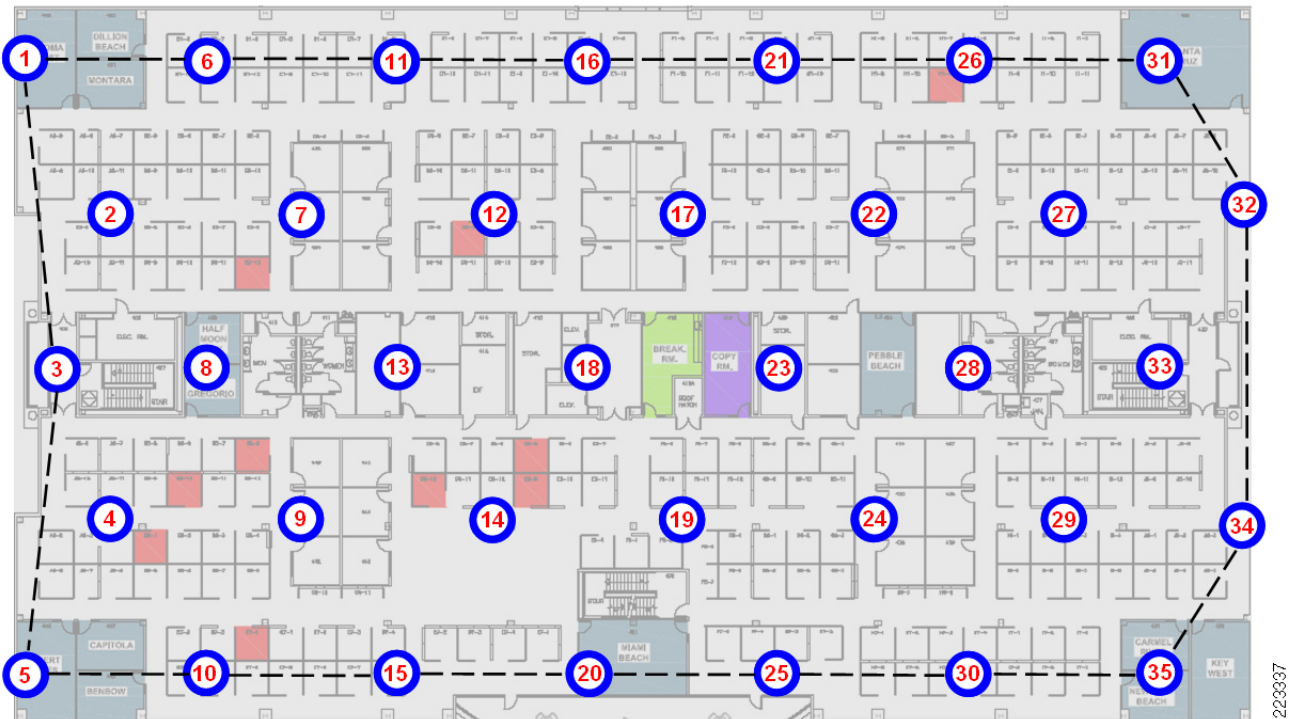
**Note**

---

Smartphones like iOS or Android devices double up both as data and voice clients due to the inherent nature of devices and applications that can be loaded on them. For the purposes of this example, we shall treat smartphones as voice devices that can also do data and follow get a voice optimized design to be location ready as well.

---

Figure D-1 Layout for 5GHz Voice and High Speed Data—2.4GHz Legacy



In Figure D-1, we assume the use of 35 ceiling mounted AP3700 access points, each of which is equipped with a pair of 2.2dBi antennas for 802.11bgn and a pair of 3.5 dBi antennas for 802.11an/ac. The access points and the antennas are mounted at a height of 10 feet. The design is intended to provide a minimum of -67 dBm signal level and a data rate of at least 24 Mbps on 802.11an for VoWLAN and high speed data clients, and a minimum of -67 dBm signal level and data rate of at least 11 Mbps on 802.11bgn for legacy data and voice clients. 802.11an VoWLAN devices are assumed to be Cisco IP phones with integrated antenna. Legacy voice and data client devices are assumed to possess nominal antenna gain of 0 dBi. Inter-access point spacing is approximately 42.7 feet and was selected to allow for a uniform distribution of access points within the floor interior and also ensure that the access point power levels required to produce our desired cell-to-cell overlap would fall within the capabilities of our client devices.

Note the following:

With the exception of access points 1, 5, 32, and 34, access points are not located directly at the floor perimeter. This is not optimal for the support of good location accuracy in all areas of the floor.

The lack of perimeter access points in the right hand corners of Layout for 5GHz Voice and High Speed Data, 2.4GHz Legacy. Because of this, there are areas in the vicinity of access points 31, 32, 34, and 35 where the location requirement for each point to lie within 70 feet of three different access points in at least three different quadrants (with an access point present in the fourth quadrant at any range) will not be satisfied.

Transmit power for each access point has been configured to +5dBm for 802.11bgn and +11 dBm for 802.11an/ac. This results in a -67 dBm cell radius of approximately 28.72 feet with a cell-to-cell overlap of 15% for 802.11an/ac VoWLAN and high speed data clients. For 802.11bgn legacy clients, it results in a -67 dBm cell radius of approximately 31 feet with a 20% cell-to-cell overlap.

**Note**

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

To facilitate optimal location tracking with this design, a few changes, additions, and adjustments will be necessary. Examining the current voice and data design and its associated parameters, the current access point spacing, antenna installation height, and placement pattern appear to be acceptable for location usage. However the lack of access points located at the actual floor perimeters and in the corners of the floor is a concern that should be addressed. This can be seen from the dashed line in [Figure D-1](#) that illustrates the convex hull established by the current perimeter of access points. Note that areas at each corner and along each upper and lower perimeter lie outside of this boundary. Although these areas may not prove to be a hindrance to some users, for the purposes of this example, our goal is to ensure optimal location accuracy in all areas of the floor, including the conference rooms in the corners of the floor and in all perimeter areas. Therefore establishing a proper floor perimeter will be our first order of business.

The first step is to implement top and bottom access point perimeters as close to the building perimeter as feasible, while attempting to maintain the uniform density of access points shown in [Figure D-1](#) to the highest degree possible. Maintaining a high degree of access point uniformity is especially beneficial to those users that depend on the Cisco Radio Resource Management (RRM) to maintain transmit power control and perform coverage hole remediation. RRM functions most effectively when the distribution of access points on a floor is as uniform as possible.

**Note**

While the recommendations show that Access Points are placed at absolute corners or touching the perimeters of an space, Access Points maybe placed a little away from the perimeter to avoid RF signal wastage outside of the perimeter. Designers are encouraged to have a RF Plan that also maximizes RF usage, as well as provide good perimeter coverage.

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

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


When considering the first option, it is necessary to examine the current inter-access point spacing and transmit power levels and estimate the increase that will be required to the inter-access point separation to place the existing outer rows of access points at the actual floor perimeters. If current access point transmit power levels are already at high levels relative to the power capabilities of our client devices, and the estimated increase to inter-access point separation appears to be large, then expanding the constellation of existing access points to accommodate perimeter placement may not be the best option. This is mainly because it may require the use of higher than desirable access point transmit power levels. In such cases, it is recommended to pursue the second option, which contracts the equilateral formations and results in shorter inter-access point separation, typically with the same or reduced access point transmit power levels.

Recall from our discussion that our transmit power levels are configured at +5dBm for 802.11bgn and +11 dBm for 802.11an/ac. To determine the new inter-access point separation that would be in effect if we were to uniformly expand the current formations (seen as equilateral triangles in Layout for 5GHz Voice and High Speed Data, 2.4GHz Legacy), we need to perform some basic geometrical calculations. We determine the new inter-access point separation required by assuming that the current top and bottom rows of access points are relocated such that they are positioned at the actual top and bottom floor perimeter. For the 275 x 159 floor in [Figure D-1](#), this is performed by dividing the top-to-bottom width of the floor (159 feet) by the number of desired rows of equilateral triangular formations (4), thereby yielding a projected formation height of 39.75 feet.

From the premise that in an equilateral triangle each angle is equal to  $60^{\circ}$  (shown in [Figure D-2](#)), we calculate the length of any side  $s$  from the height  $h$  of our equilateral triangle formations as follows:

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

**Figure D-2** Equilateral Access Point Formation



Solving for  $s$ , we calculate:

$$s = \frac{h}{\sin 60^{\circ}}$$

Or:

$$\frac{39.75}{.866}$$

= 45.9 feet. Thus we would need to expand our current inter-access point spacing from 42.7 feet to 45.9 feet to move both the top and bottom rows of outermost access points to the actual building perimeter. As this represents a relatively minor increase in inter-access point spacing, it should be easily

accommodated by a correspondingly minor increase in transmit power, if any at all. In our next step, we determine the new cell size that would be required to support the recommended levels of overlap, given our newly calculated inter-access point spacing.

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

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

Where:

- $O$  = the desired overlap percentage divided by 100

$$\arccos\left(\frac{d}{2R}\right)$$

is expressed in radians

- $d$  = the inter-access point distance in feet
- $R$  = the cell radius in feet

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

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

$$TX_{POWER} = RX_{POWER} + Loss_{TX} - Gain_{TX} + PL_{REFERENCE} + 10 \log D^N + s - Gain_{RX} + Loss_{RX}$$

For the purposes of this example, we have assumed:

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

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

**802.11bgn:**

$$\begin{aligned} TX_{POWER} &= -67 \text{ dBm} + 0 - 2.2 \text{ dBi} + 40 \text{ dB} + 10\log(10.183.5) - 0 + 0 \\ &= -29.2 + (10 * 3.527) \end{aligned}$$

$$TX_{POWER} = +6.07 \text{ dBm, or approximately } +8 \text{ dBm}$$

**802.11an/ac:**

$$\begin{aligned} TX_{POWER} &= -67 \text{ dBm} + 0 - 3.5 \text{ dBi} + 46 \text{ dB} + 10\log(9.413.5) - (-3.0) + 0 \\ &= -24.5 + (10 * 3.408) + 3 \end{aligned}$$

$$TX_{POWER} = +12.58 \text{ dBm, or approximately } +14 \text{ dBm}$$

Note that these power levels have been rounded upward to the next available transmit power increment available on the AP3700 access point. Since this is +1.93 dBm higher than the required transmit power to achieve our recommended 20% overlap goal at a cell signal boundary of -67 dBm, we can expect that the overlap will exceed the 20% target. This is acceptable, as the 20% overlap is a minimum target. Similarly, for 802.11an/ac the access point transmit power level of +14 dBm is +1.42 dBm higher than what is required to achieve the recommended 15% overlap, once again resulting in more overlap between cells than expected.

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

**Note**


---

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

---

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

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



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

Figure D-3 illustrates the updated access point layout using the information from the calculations above along with perimeter access point placement which is discussed next.

**Figure D-3** Layout for 5GHz Voice and High Speed Data—2.4GHz Legacy with Location



In Figure D-3 we can see the effects of the increase in inter-access point distance:

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

The two new access points added in [Figure D-3](#) bring the total access point count for the integrated voice, data, and location design to 37 access points. The primary source of voice and data coverage in this design still emanates from the access points participating in the equilateral formations seen across the floor (i.e., this can be seen in [Figure D-3](#) as the set of access points depicted in red). Access points 32, 34, 36, and 37 are necessary to establish a location perimeter, but based on the assumptions and calculations presented here, may not be required to participate in providing voice or data coverage in either band. That being the case, these access points can be statically configured to operate at significantly reduced transmit power (such as -1 dBm, for example), which also minimizes the co-channel interference contribution of these access points as well.

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

Based on our planning output and our calculations, our original voice and data design shown in [Figure D-1](#) can be migrated to a location-ready design. The result is a combined design that is well suited to support VoWLAN, high speed data and location tracking on 5 GHz, as well as legacy data and voice support with location tracking on 2.4 GHz.

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



## CMX Troubleshooting

---

Revised: September 4, 2014

The appendix lists common deployment issues and their solutions.

### MSE and WLC Communication Problems

Sometimes with a Virtual install of Mobility Services Engine, there can be a misconfiguration of certificate between WLC and MSE. If the NMSP status between the MSE and WLC is not in an “UP” state, verify that the following steps have been followed:

1. MSE and WLC are time synced to the same NTP server.
2. WLC time should not be behind MSE time. Also any time difference between the two should not be greater than 1-2 minutes.

If after verifying these steps the MSE and WLC NMSP connection is still not up, then it is important to:

- Configure MSE manually on the WLC.
- Use un-authenticated NMSP tunnel between MSE and WLC.

Perform the following.

---

**Step 1** Login into MSE via ssh connection or via a console connection.

**Step 2** Issue the commands:

```
[root@cmxmse ~]# cmdshell -- Issue this command to get into cmdshell

cmd> show server-auth-info - Issue this command to get auth info for the MSE
invoke command: com.aes.server.cli.CmdGetServerAuthInfo
AesLog queue high mark: 50000
AesLog queue low mark: 500
-----
Server Auth Info
-----
MAC Address: 00:0c:29:b1:f5:a8 - Note the MAC address
SHA1 Key Hash: ee68b5062b4181f68d5dd489db2bfcf5637b5eff
SHA2 Key Hash: ec7ebc55bbe366332da70e995f2c073bc7cfaf4cb6d845336adfc67ce961644 -- Make a
note of this key to be used later
Certificate Type: SSC

cmd> config unauthenticated-nmsp true - Enable Un-authenticated NSMP connection.
invoke command: com.aes.server.cli.CmdSetServerConfigParameter
```

Parameter unauthenticated-nmsp was successfully modified

**Step 3** Login into the WLC via SSH or Console shell and invoke the commands:

```
(Cisco Controller) >config auth-list add sha256-lbs-ssc <MAC ADDRESS> <KEY HASH>
MAC ADDRESS and KEY HASH are derived from Step 2.
```

**Step 4** Verify that MSE has been manually added on the WLC and the NSMP connections are up between the two. Invoke the following commands on WLC:

```
(Cisco Controller) >show auth-list - Shows the manually added MSE to WLC
```

```
Authorize MIC APs against Auth-list or AAA ..... disabled
Authorize LSC APs against Auth-List ..... disabled
APs Allowed to Join
  AP with Manufacturing Installed Certificate.... yes
  AP with Self-Signed Certificate..... yes
  AP with Locally Significant Certificate..... yes
```

Mac Addr	Cert Type	Key Hash
00:0c:29:b1:f5:a8	LBS-SSC-SHA256	ec7ebc55bbef366332da70e995f2c073bc7cfaf4cb6d845336adfc67ce961644

```
(Cisco Controller) >show nmsp status - Shows NSMP status
```

MSE IP Address	Tx Echo Resp	Rx Echo Req	Tx Data	Rx Data
<MSE IP >	75779	75779	210547	12

## Aspect Ratio Issues while Creating Maps

When uploading maps into Cisco Prime Infrastructure, ensure that the scale of the map is correct. If the scale is not correct, there will be aspect ratio issues on the map. If you see that the aspect ratio of the map you upload is not correct, delete the map and re-upload it with the correct scale. There is no way to modify scale after the map has been uploaded.

## Coverage Zones Cannot Be Renamed

Ensure that Coverage Zones have the right names before deploying the solutions. The names of Coverage Zones cannot be renamed in this release.