Real-Time Traffic over Wireless LAN Solution Reference Network Design Guide

First Published: November 11, 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-29731-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

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.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE 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 THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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 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)

© 2013 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface

This preface describes the purpose, audience, and document organization.

Purpose

The *Real-Time Traffic over Wireless LAN Solution Reference Network Design Guide* provides a design reference for wireless solutions that provide connectivity for endpoints and clients that send and receive real-time traffic, and utilize real-time traffic applications and services. Wireless network deployments that support real-time traffic capable endpoints and enable real-time traffic applications and services are referred to as Real-Time Traffic over WLAN (RToWLAN) deployments.

Real-time traffic endpoints and applications generate and consume real-time network traffic. This network traffic includes packetized voice and video, as well as other traffic consumed as near to the moment it is generated as possible. Because the value of real-time network traffic drops to zero almost instantly, there is no retransmission, and limited tolerance for delay, variation in delay (jitter), or packet loss. The network must deliver real-time traffic between transmitters and receivers with negligible delay and packet loss; otherwise retransmission and delayed traffic may traverse the network only to be discarded at the far end by the receiver.

A well-planned RToWLAN deployment design not only provides high-quality voice and video communications but also provides sufficient delivery times for other real-time traffic applications and services like desktop virtualization and presence. This design guide focuses on solution-level planning-and design-related aspects of RToWLAN deployments rather than on specific hardware and software requirements.

The *Real-Time Traffic over Wireless LAN Solution Reference Network Design Guide* supersedes and deprecates the previous *Voice over Wireless LAN Design Guide* that is available at http://www.cisco.com/.

Audience

This guide is intended for systems design and deployment engineers who are responsible for planning and designing the Cisco Unified Wireless LAN deployments for real-time traffic endpoints and clients.

Organization

The following table lists the chapters of this guide:

Table 1: Guide Overview

Chapter	Description
Preface, on page ?	Describes the purpose, audience, and document organization.
Real-time Traffic over WLAN Overview, on page 1	Provides an overview of RToWLAN solution, the solution architecture, high-level design information related to wireless, collaboration, endpoints, and network management including the various components and considerations for RToWLAN deployments.
Real-Time Traffic over WLAN Radio Frequency Design, on page 21	Provides an overview of the radio frequency (RF) network requirements of RToWLAN deployments and RF deployment issues.
Real-Time Traffic over WLAN Quality of Service, on page 47	Provides an overview of WLAN QoS and its implementation in the Cisco Unified Wireless Network.
Real-time Traffic over WLAN Security, on page 79	Provides an overview of WLAN security.
Real-time Traffic over WLAN Roaming, on page 91	Provides an overview of WLAN roaming and implications for RToWLAN deployments.

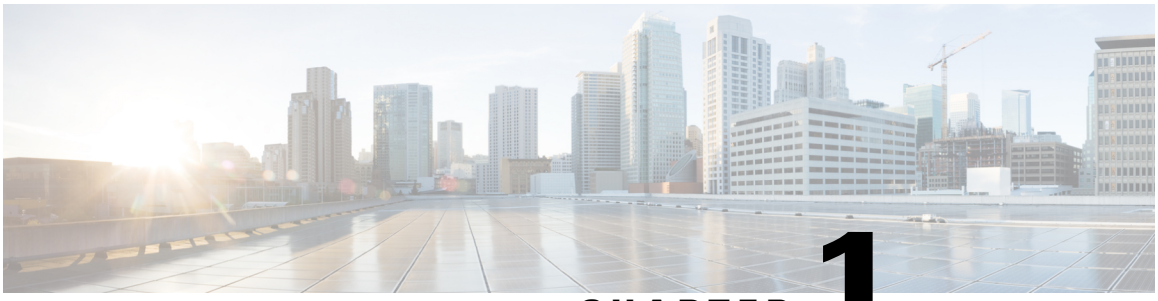

CHAPTER 1

Real-time Traffic over WLAN Overview

This chapter discusses about drivers and benefits of implementing an RToWLAN deployment followed by an illustration that depicts an enterprise solution reference network architecture. This chapter also identifies the high-level touch points of the RToWLAN solution deployment, provides a brief overview of RToWLAN solution architecture, and describes the following three main RToWLAN solution components:

- 802.11 enterprise WLAN solution infrastructure
- Enterprise collaboration solution applications and services
- Real-time Traffic over WLAN endpoints

After the description of the components, the chapter describes common RToWLAN solution design considerations across both single-site and distributed multisite deployments with specific focus on quality of service (QoS), security, high availability, and capacity planning.

RToWLAN solution drivers and benefits

Enterprises today are faster-paced than ever before. To succeed, grow, and stay ahead of the competition, enterprises depend on efficient employees, collaboration, and a timely business process. While maintaining a secure corporate data and communications infrastructure, enterprises are encouraging a mobile work style to get more work done. Additionally, enterprises are not just streamlining processes through technology but

are also looking to increase their revenue and reduce costs with new technology that improves user productivity and accelerates business processes.

Real-time traffic applications and services that are delivered over WLAN networks provide the following benefits:

- **Eliminates the need for mobile cellular devices in the enterprise:** IP voice and video over WLAN calls traverse the enterprise WLAN in whole or in part, providing cost savings over direct cellular network calls. Employees can use their voice or video over WLAN endpoint or client on campus instead of consuming voice minutes over the cellular voice network.
- **Reduces dependency on mobile provider network coverage within the enterprise:** By leveraging 802.11 WLAN network connectivity, enterprises provide adequate network coverage and capacity by deploying sufficient number of access points and at the same time reducing or eliminating dependency on mobile provider network coverage within the enterprise.
- **Enablement of employee-owned personal and guest devices:** With the prevalence of personal mobile devices like smartphones and tablets, there is an increasing inflow of these devices into the enterprise. This type of enterprise is often called *bring your own device* (BYOD). Enabling employee-owned or guest devices with a BYOD solution:
  - increases overall employee satisfaction.
  - improves productivity when enabling device for real-time traffic collaboration applications and services.
- **Maximizes availability and reachability of mobile employees:** By enabling mobile devices for enterprise collaboration, you can reach employees anywhere within the enterprise, and the user experience on the mobile device is equivalent to that of the traditional enterprise endpoint. This flexibility provides a smooth experience to users who transition between many enterprise clients and devices.
- **Ensures high-quality voice and video calls and seamless mobile user experience for fixed mobile substitution (FMS) deployments:** To enable enterprise directory number use on dual-mode mobile smartphones and tablets through enterprise IP telephony systems, you must send IP voice and video calls over the enterprise WLAN network. A RToWLAN network is tuned to provide optimal bandwidth and throughput as well as higher-priority queuing or transmission on the network for real-time traffic. In the case of IP-based voice and video calls, this means minimized packet loss, jitter, and delay translating to high-quality voice and video. Enabling the enterprise directory number usage on mobile endpoints and clients:
  - provides integration to enterprise IP call-control features.
  - provides enterprise dial plan, enabling seamless user experience across enterprise devices.
  - eliminates the need for an enterprise deskphone for mobile employees.

RToWLAN deployments can help enterprises to:

- Reduce expenditures for mobile provider voice and data services.
- Improve employee productivity, reachability, and availability.
- Leverage the increased presence of personal mobile devices within the enterprise for collaboration and communication at less or no cost with BYOD solutions.
- Improve employee satisfaction when they engage with collaboration and other business applications and services by providing flexibility and a seamless user experience.

RToWLAN solution reference network design architecture

This section provides a high-level example network topology (see Figure 1: Real-Time Traffic over WLAN Solution Network Topology Overview, on page 3) for deploying real-time traffic applications and services for wireless endpoints and clients. This RToWLAN solution design example uses a typical hierarchical, access, distribution, and core campus network as a basis. The following additional components are included in this design:

- WLAN Controller (WLC) and Access Points (APs) to provide the wireless network infrastructure for carrying real-time and other IP network traffic.
- Wireless access points added to Client and Endpoint Access Blocks for client and mobile endpoint wireless network attachment.
- Data Center Block including call control and other application servers for enabling real-time traffic.
- Voice and Internet Gateway Block for access to PSTN and to provide access to and from the enterprise.

Figure 1: Real-Time Traffic over WLAN Solution Network Topology Overview

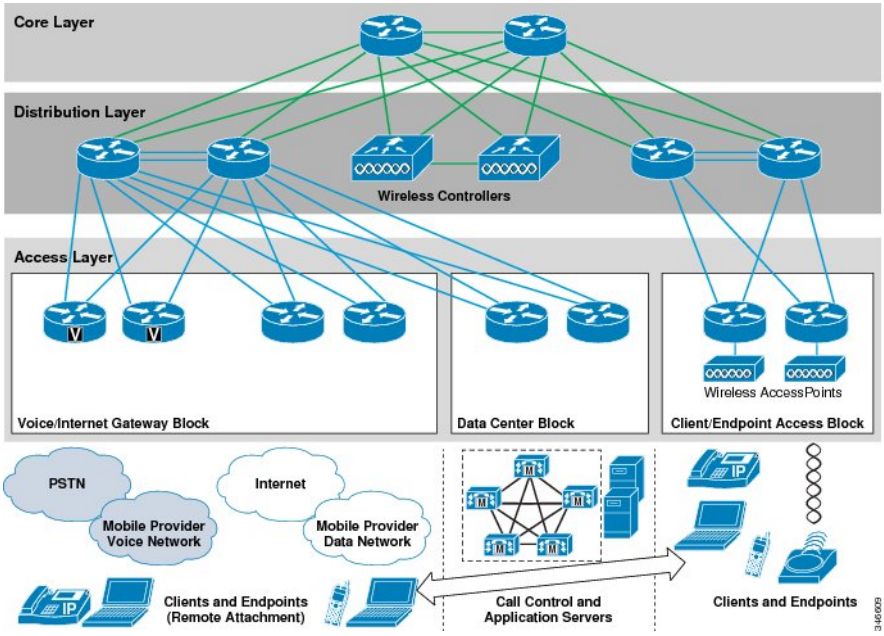

The two key solution component areas of focus for any enterprise RToWLAN deployment are:

- Enterprise 802.11 wireless infrastructure enabled within the Distribution and Client and Endpoint Access blocks as shown in the preceding figure.
- Enterprise collaboration infrastructure enabled within the Data Center and Voice and Internet gateway as shown in the preceding figure.

Enterprise wireless LAN overview

The enterprise 802.11 wireless LAN (WLAN) infrastructure is the underlying foundation for any RToWLAN solution deployment. The enterprise wireless network must be designed to ensure that network connectivity is available for wireless endpoints and that sufficient bandwidth and throughput for real-time traffic is provided. The WLAN must be designed with sufficient capacity for the number of RToWLAN endpoint devices expected to be deployed. The WLAN must also be designed with sufficient redundancy such that hardware or IP connection failures do not completely eliminate WLAN network availability.

The WLAN infrastructure not only provides network connectivity but just as with wired networks, the infrastructure also provides authentication and encryption security services and quality of service for better than best effort treatment for select traffic. To deliver this functionality, the 802.11 WLAN infrastructure consists of a number of components and applications including wireless LAN controllers (WLCs), wireless access points (APs), and wireless LAN management applications.

Enterprise collaboration overview

Collaboration systems enable a large number of features and services. The most common and prevalent feature is voice and video over IP calling. However, these collaboration systems can provide communication features and functions above and beyond traditional IP telephony including conferencing, messaging, presence, information and document sharing, fixed mobile convergence, and directory integration. These features and services are often deployed in tandem, providing a comprehensive collaboration solution for enterprises and their workers. To deliver these features and services, the collaboration system relies on a number of components and applications including voice and video endpoints, gateways and application servers including voicemail and presence.

RToWLAN solution architecture overview

The overall architecture for an RToWLAN solution deployment (see Figure 2: RToWLAN solution architecture overview, on page 5) consists of the following three main components:

- **802.11 Wireless LAN infrastructure:** The wireless infrastructure enables the 802.11 wireless LAN for endpoint or client attachment. This infrastructure includes the WLAN controller, access points, and management applications.
- **Collaboration applications and services:** Collaboration applications enable real-time traffic services, including voice and video calling. These applications and services include call control, PSTN gateways, media resources, voicemail, and instant messaging and presence.
- **Real-time traffic capable 802.11 wireless endpoints:** Wireless endpoints consume and generate real-time traffic over the 802.11 WLAN. These endpoints include wireless-enabled desktop and mobile software applications and clients as well as wireless IP phone hardware.

All of these components are involved in enabling RToWLAN applications and delivering real-time services wirelessly.

Figure 2: RToWLAN solution architecture overview

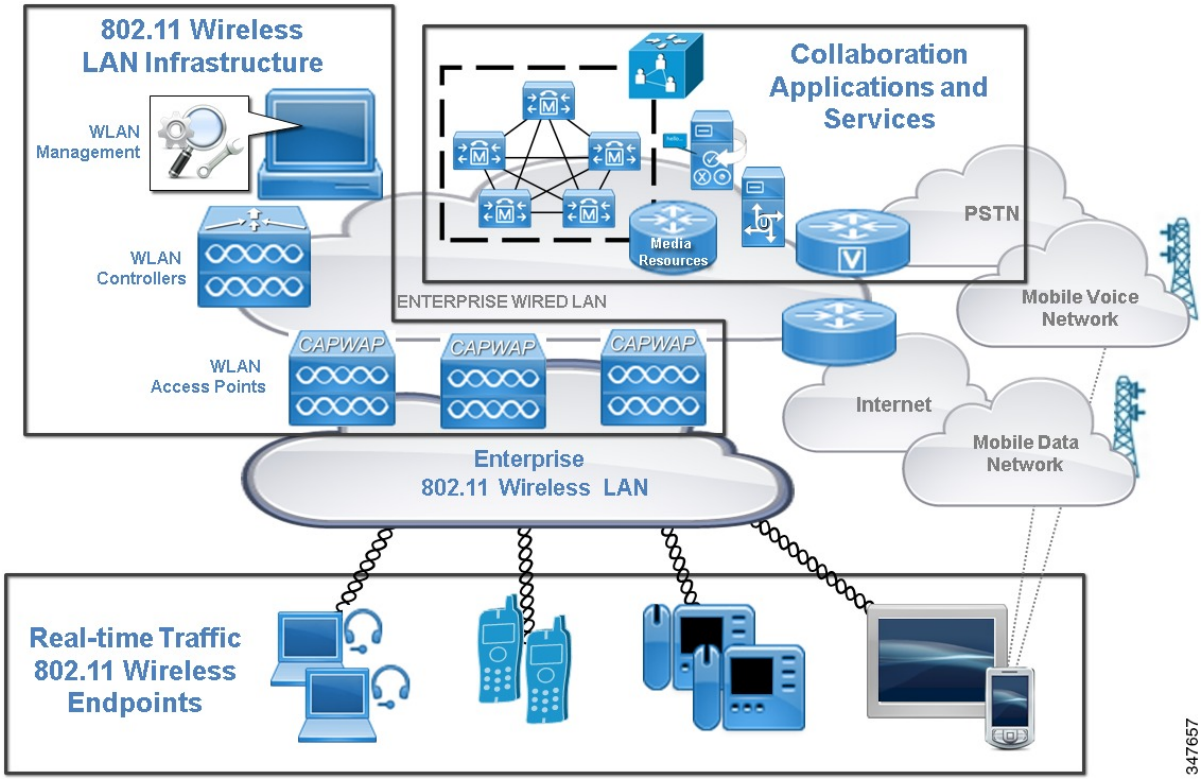

See Figure 2: RToWLAN solution architecture overview, on page 5 above and subsequent discussions within this chapter do not include information about basic network services that are provided by typical enterprise applications and servers. These network services are assumed to be present, and considerations related to the following are not discussed except in cases that are directly related to RToWLAN:

- Network-based device and user authentication and identification services, including certificate authority servers, two-factor authentication supplicants, identity stores like directory servers, and any other applications or components that provide security services.
- Network time and IP address resolution and assignment, including network time services (NTP), domain name services (DNS), and dynamic IP address assignment (DHCP).
- Network routing, packet forwarding and queuing, quality of service, and admission control.

Enterprise 802.11 wireless LAN solution infrastructure

The enterprise 802.11 wireless LAN (WLAN) network is essential for RToWLAN solutions, because the wireless infrastructure provides the network that transports real-time traffic that connected real-time traffic capable wireless devices generate and consume. The following table lists the 802.11 wireless network infrastructure components:

Table 2: 802.11 Wireless Network Infrastructure Components

802.11 wireless network infrastructure components	Description
Wireless LAN access points	The wireless LAN access point provides wireless network access to wireless devices that enables the devices and clients to communicate with wired network components. The access points not only provide wireless device network connectivity, but they also serve as the demarcation point between the wired and wireless networks. The wireless LAN controllers manage the access points that are registered to them.
Wireless LAN controllers	The wireless LAN controller (WLC) is a network infrastructure device that performs the central management role in the wireless network. WLCs make it easier to manage wireless LAN deployments by centralizing access point configuration and management, radio frequency monitoring, and client association and authentication. After the wireless access points register to a WLC, the wireless access points tunnel all management and client traffic to the WLC. The WLC is also responsible for switching traffic between wireless clients and the wired portion of the network.
Wireless management	Wireless LAN management applications and services provide a robust wireless life-cycle management tool that enables network administrators to successfully plan, deploy, monitor, troubleshoot, and report on wireless networks.

You must design the enterprise WLAN network to meet the needs of the users, applications, and endpoints. The WLAN coverage must provide enough bandwidth on the Wi-Fi channels to support quality application performance. To design enough bandwidth into WLAN coverage areas for the users, you must understand the Wi-Fi performance capabilities of the endpoints. Wireless endpoints and mobile client devices are available in different varieties, and not all wireless clients have the same capabilities. You must perform detailed wireless infrastructure planning for the WLAN deployment to be successful.

For a successful wireless network deployment, you must conduct a thorough wireless site survey to ensure that the radio frequency configuration and design are optimized to provide necessary bandwidth and throughput to endpoints in motion throughout the deployment. In addition, the site survey helps you identify the sources of interference so that they can be eliminated. A site survey seeks to verify the following basic radio frequency design principles:

- **Adjacent channel cell separation:** Same channel or adjacent channel can cause interference, which reduces network throughput and leads to increased packet loss. Through site survey, you can verify that adjacent channel cells are separated appropriately.
- **Nonadjacent channel cell overlap:** Nonadjacent channel cells should be overlapped to ensure that wireless endpoints can seamlessly transition or roam between access points and wireless channel cells. Through site survey, you can verify that nonadjacent channel cells are overlapped sufficiently.
- **Channel cell coverage for all desired areas of service:** If wireless network coverage is expected in stairwells, between buildings, or on building perimeters, you can verify through site survey that access

points and antennas are located properly to provide appropriate wireless channel coverage in all required locations.

- **Channel cell density:** Appropriate channel cell density must be provided such that the required number of endpoints are supported with necessary network bandwidth and throughput provided on WLAN channels to support quality real-time application and service performance.
- **Wireless interference identification and mitigation:** Wireless interference caused by improper AP and antenna placement, physical structures and characteristics of the deployment area, and poor radio frequency design must be identified during a site survey. Further, sources of interference such as cordless phones, personal wireless network devices, sulphur plasma lighting systems, microwave ovens as well as high-power electrical devices such as transformers, heavy-duty electric motors, refrigerators, elevators and elevator equipment, and any other power devices that could cause electromagnetic interference (EMI) must be considered when you plan a WLAN deployments. You must mitigate these interference sources by adjusting access point locations and antenna direction, radio frequency configuration and AP power levels, or by removing or eliminating these interference sources.

Today, controllers and access points are purpose-built for particular use cases and levels of scale. In a home environment, a limited number of mobile client devices share the same WLAN channel radio frequencies at the same time. Hence, bandwidth is not a concern. However, in an enterprise environment, many devices are connected to a Wi-Fi channel. With inadequate bandwidth, voice calls may become inaudible, video calls may become unwatchable, and in some cases, the application fails.

Related Topics

Real-Time Traffic over WLAN Radio Frequency Design, on page 21
Cisco wireless products
Enterprise Mobility Design Guide

Enterprise collaboration solution applications and services

After the enterprise wired and wireless LAN is planned and deployed, you must deploy the enterprise collaboration applications, services, and endpoints on top of this infrastructure. The enterprise collaboration deployment must be properly designed to ensure that required components, applications, and services are available and that the infrastructure provides sufficient capacity and reliability.

Collaboration systems enable a large number of features and services, including voice and video calling; messaging, including voicemail and instant messaging (IM), presence and availability; media resources, including conferencing and music on hold, and directory integration. These features and services are often deployed in tandem, which provides a comprehensive collaboration solution for enterprises and their workers.

The enterprise call control platform, which is considered central to the collaboration system, is responsible for providing voice and video calling services (for example, Cisco Unified Communications Manager). The collaboration system also relies on several components and applications, including the following:

- **Gateways:** Platforms with IP or TDM interfaces that provide external access to the PSTN as well as internal access to other call control platforms, applications, and devices.
- **Media Resources:** Hardware or software-based resources that are deployed throughout the network to enhance call flows with supplementary services like voice and video conferencing, music on hold, and transcoding.

- **Collaboration applications:** Applications that provide communication features and functions beyond traditional voice and video calling. These include voicemail, IM and presence, conferencing, information and document sharing, fixed mobile convergence, and directory integration.

Voice and video endpoints are also a key component of the collaboration system. Enterprise users communicate and collaborate with endpoints including desk phones, wireless phones, software clients, immersive video systems, and mobile clients for smartphones and tablets.

When deploying collaboration hardware, applications and services, and before enabling real-time traffic capable wireless endpoints, ensure that the required collaboration and communication applications and services are deployed and configured appropriately. You must consider the following factors:

- Device configuration

  Device configuration within the call control system including, location, user association, calling privileges, button layout, and other feature settings must be configured correctly for proper operation.

- Network service operations

  Devices must be capable of attaching to the network, retrieving or receiving their configuration from the call control system over the network, and connecting to and communicating with the collaboration infrastructure components.

- Dial plan

  The enterprise dial plan that is configured on the call control system is as fundamental to collaboration system operation as IP addressing and routing are to IP networking. You must implement the following dial plan configuration correctly to ensure that users are capable of making and receiving calls:

  - The numbering or assigning of directory or telephone numbers to devices
  - The dialing habits enabled and the classes of restriction applied to devices
  - The manipulation of called or calling numbers as configured for the device

- Intercomponent integration

  Enable and configure integrations between collaboration and communication infrastructure components to deliver additional features and services. For example, integrations between the call control platform and PSTN gateways or border controllers, the voicemail and IM or presence application servers, and directory or contact source platform must be put in place to enable additional services like PSTN access or IM and presence.

- Security

  Deploy security features as indicated by enterprise security policies to ensure that operations are secured. Specifically, if you plan to use functionality such as encrypted voice or video media or signaling, certificate-based trunk integrations, or digest authenticated devices, then you must deploy an appropriate configuration and infrastructure, including a certification authority server.

Use a nonwireless real-time traffic capable endpoint first to verify the preceding considerations and ensure that the required features and functions are working correctly before proceeding with deployment of RToWLAN wireless endpoints.

Related Topics

802.11 RToWLAN endpoints, on page 9
Cisco voice products

Cisco collaboration products
Cisco enterprise collaboration deployment design

802.11 RToWLAN endpoints

After you design and implement the enterprise WLAN infrastructure and collaboration applications and services, you must deploy the RToWLAN-capable devices and clients on top of this overall infrastructure.

Note Consider the RToWLAN endpoint selection during the WLAN infrastructure design and implementation phase, particularly during the site survey. If you design the WLAN infrastructure without consideration for the specific RToWLAN endpoints that will be deployed on this infrastructure, the result will be issues such as packet loss, excessive delay, and poor voice and video quality.

RToWLAN endpoints are categorized into two categories:

- Hardware-based wireless IP phones
- Software-based clients for wireless devices

Hardware-based wireless IP phones

The hardware-based wireless IP phones are purpose-built wireless IP voice and video phones that are designed specifically to provide voice and video calling functionality.

Hardware-based wireless IP phones fall into two subcategories:

- **Desktop WLAN Phones (or Fixed WLAN Phones):** IP desk phones connect to the network wirelessly but remain stationary because they are powered from a wall outlet. For example, Cisco Desktop Collaboration Experience DX650.
- **Mobile WLAN Phones:** Wireless IP phones are battery powered and are connected to the network wirelessly, which enables them to move throughout the enterprise while maintaining an active network connection. For example, Cisco Unified Wireless IP Phone 7925G.

Software-based clients for wireless devices

The software-based clients for wireless devices run on multifunction wireless devices that are capable of generating many types of traffic and performing many types of operations. The multifunction wireless device runs a collaboration software application that enables voice and video calling as well as other collaboration features.

Software-based clients running on wireless devices fall into two subcategories:

- Desktop computing platforms

  These devices connect to the network wirelessly, but depending on the platform, they may or may not remain stationary. Desktop computers rely on power from a wall outlet, while laptop computers are battery powered, and can therefore move around while maintaining an active wireless connection to the network. These devices typically run Microsoft Windows or Apple Mac operating systems, so the collaboration software client must support these operating systems.

  For example, Cisco Jabber for Windows.

- Mobile computing platforms

  Smartphones and tablets are examples of these types of devices. They are battery powered, and when wirelessly connected, they can move throughout the enterprise while actively maintaining connection to the network. These devices typically run Android or Apple iOS operating systems, so the collaboration software client must support these operating systems.

  For example, Cisco Jabber for Android.

Endpoint selection and WLAN site survey

You must understand the WLAN performance capabilities on the endpoints to design enough bandwidth into WLAN coverage areas for the users. Wireless endpoints and mobile client devices are available in different varieties, and not all wireless clients have the same capabilities.

Site surveys are one of the basic requirements when you deploy a WLAN, and you must always consider the Wi-Fi capabilities of the client devices or endpoints. Most smartphones and tablets support 802.11. However, generally, the smartphones and tablets have fewer antennas and lower data rates than laptops. In addition, most are not purpose-built for the enterprise WLAN market. Almost all smartphones and tablets support enterprise security policies. However, many of them do not support Wi-Fi protocols and features like 802.11e WMM for QoS and bandwidth control for audio and video calls. Consider these limitations before you start the survey process.

Smartphones and tablets generally have subpar access point to access point roaming logic. Most consumer or nonenterprise Wi-Fi endpoints perform poorly when they roam between access points. Client roaming can be unpredictable, and it is common for an endpoint without enterprise roaming logic to travel the same path repeatedly without repeating the same roam times in milliseconds or seconds. Nonenterprise Wi-Fi endpoints often repeatedly roam to the same access points in a high-density deployment rather than using other access points that may provide better throughput. It is important to test and understand the capabilities of the devices that are to be used in your facilities before you start the survey process.

Cisco WLAN controllers provide parameters to help nonenterprise clients roam. These parameters include minimum received signal strength indication (RSSI), hysteresis of decibels (dB), scan threshold in decibels per milliwatt (dBm), and transition time. The settings of these parameters need to be tested on-site with different mobile clients, because mobile clients behave differently.

Related Topics

Enterprise collaboration solution applications and services, on page 7
Real-Time Traffic over WLAN Radio Frequency Design, on page 21
Cisco IP phones
Cisco Jabber

RToWLAN solution deployment considerations

This section describes important factors that you must consider when you design and implement an RToWLAN solution.

Hardware and software selection

When you select the hardware and software for an RToWLAN solution deployment, you must consider feature sets, standards and capabilities that are supported, and compatibility of the hardware and software. It is important to ensure that the selected wireless and collaboration infrastructure and the devices that are deployed

on that infrastructure deliver the required features and capabilities, whether you are selecting wireless LAN controllers and access points, collaboration platforms and applications, or the endpoint devices.

As a general rule, you should choose products that support a rich set of advanced network and application features, and at the same time, meet approved wireless and collaboration protocol standards to ensure interoperability and compatibility with a wide variety of systems. For example, when you choose wireless infrastructure components, consider full support for 802.11 wireless standards (802.11a, 802.11g, and the newer 802.11n and 802.11ac wireless access standards, as well as advanced wireless standards like 802.11e and 802.11r) as minimal requirements. Support for these standards ensures that necessary bandwidth with minimal delay and best effort treatment are provided for real-time traffic like voice and video.

When you design collaboration and communications infrastructure, a system that provides advanced capabilities (location and availability awareness, fixed mobile convergence, voice and video over IP, dual-mode device support, and so on) is needed to deliver the appropriate feature set that mobile workers require with RToWLAN deployments.

Voice and video over WLAN

It is critical that you plan and deploy a finely tuned, QoS-enabled, and highly available WLAN network to enable voice and video calling and other real-time traffic applications to ensure a successful RToWLAN solution deployment.

Because the 802.11 RToWLAN endpoints rely on the WLAN infrastructure for carrying both critical call signaling and real-time voice and video media traffic, you must deploy a WLAN network that is optimized for both data and real-time media traffic. A poorly deployed WLAN network results in a large amount of interference and diminished capacity, leading to poor RToWLAN application and service performance. In the case of voice and video calling, issues include not only poor call quality but in some cases dropped or missed calls. The poor application performance renders the WLAN deployment unusable for making and receiving calls or using other real-time applications.

Another basic requirement is that you must conduct a WLAN radio frequency (RF) site survey before, during, and after deploying an RToWLAN solution. This ensures that cell boundaries, configuration and feature settings, capacity, and redundancy are optimized to support RToWLAN applications and services. The site survey must verify that the WLAN RF design minimizes same-channel interference and also provides sufficient radio signal levels and nonadjacent channel overlap which helps to maintain acceptable real-time traffic throughput and voice and video quality as the RToWLAN endpoint device moves or roams from one location to another.

With appropriate site survey and careful planning, the wireless infrastructure conforms to the following collaboration and unified communications application minimum network requirements:

- Average IP packet loss for collaboration or other communications application traffic of less than or equal to one percent.
- Average end-to-end delay variation or jitter for collaboration or other communications application traffic of less than or equal to 30 ms.
- Average one-way packet delay for collaboration or other communications application traffic of less than or equal to 150 ms.

Note If you implement an RToWLAN network that is intended to carry voice or video traffic where the one-way delay exceeds 150 ms, it introduces issues not only with the quality of the voice and video calls but also with call setup and media cut-through times. These problems occur because several call signaling messages must be exchanged between each endpoint and the call control platform to establish the call.

While conducting a site survey and carefully planning an RToWLAN network ensures a successful deployment on the 2.4 GHz WLAN band (802.11b/g/n), Cisco recommends that you use the 5 GHz WLAN band (802.11a/n/ac) whenever possible for RToWLAN endpoint connectivity. 5 GHz WLANs enable higher density device deployments and provide better traffic throughput and less interference than 2.4 GHz WLANs. Higher density, higher throughput, and less interference are important network characteristics for RToWLAN applications and services, that include voice and video calling. In addition, with the prevalence of Bluetooth headsets and other Bluetooth peripherals, interference on enterprise 2.4 GHz WLANs is hard to avoid. When you use the 5 GHz band for RToWLAN deployments, Bluetooth interference is not a concern.

Note In dual-band WLANs (WLANs with both 2.4 GHz and 5 GHz bands), devices can roam between 802.11b/g/n and 802.11a/n with the same service set identifier (SSID), provided the RToWLAN endpoint is capable of supporting both bands. However, with some devices, a dual-band WLAN can cause gaps in the real-time traffic path. To avoid these gaps, use only one band for real-time traffic applications and services.

Quality of Service

One critical component for successful RToWLAN solution deployments is to implement Quality of Service (QoS) at the network and application layer. QoS ensures that different types of network traffic are given access to specific amounts of bandwidth or are given priority over other traffic as they traverse the network. You can use a variety of methods to provide different levels of network throughput and access based on traffic type.

For real-time traffic, QoS methods fall into the following two categories:

- Packet marking

  Packet marking determines how packets are queued as they ingress and egress network interfaces along the traffic path. Based on packet marking, certain types of traffic are allocated more or less bandwidth or can be transmitted more quickly and more often. Generally, when traversing the network, real-time media traffic is given priority treatment in all transmit queues along the network path. Real-time signaling traffic that is used to set up calls or facilitate application features is allocated dedicated bandwidth amounts based on the expected overhead of this signaling and other control plane traffic. Real-time signaling and other non-media traffic must never be assigned to priority traffic queues.

- Packet queuing

  Packet marking may or may not be performed at the application or endpoint level, but most IP networks are capable of marking or re-marking traffic flows as they traverse the network. Marking or re-marking of traffic flows by the network is usually based on IP port numbers or IP addresses. The client application or device performs the packet marking at the endpoint level based on specific application requirements or based on standardized marking guidelines (for example, voice media should be marked as Layer 2 802.11 WLAN User Priority 6, and Layer 3 IP packet marking should be Class of Service of 5, Differentiated Services Code Point of 0x46, or Per-Hop Behavior Expedited Forwarding).

802.11 WLAN packet marking at layer 2 (User Priority, or UP) presents challenges for many RToWLAN applications and endpoints. While some applications and endpoints do mark RToWLAN traffic flows at Layer 2 according to standard guidelines, many endpoint devices, particularly multifunction mobile devices, may not support Layer 2 802.11 UP marking. Unless endpoint devices are fully 802.11e and WMM compliant, and the operating system supports UP values as marked by applications, you cannot rely on Layer 2 QoS marking to provide improved RToWLAN traffic throughput on the wireless network.

Packet marking at Layer 3 is more common in RToWLAN applications and endpoints. Many applications and endpoints mark RToWLAN traffic flows at Layer 3. When the application and endpoints are marking traffic according to recommended guidelines, existing wired network QoS policy should not be modified, because real-time traffic automatically receives appropriate treatment based on standard QoS policy (priority treatment for voice and dedicated bandwidth for video and control plane traffic).

While correct packet marking is important, whether by application or endpoint, it is also important that you trust that the correct packet marking is applied to the correct type of traffic by the application or endpoint. If the packet marking of even some network traffic generated by an RToWLAN endpoint cannot be trusted, then administrators can decide to rely on network-based packet re-marking for all traffic. In this case, all traffic is re-marked according to enterprise policy based on traffic type (port number or protocol) and IP address to ensure that network priority queuing and dedicated bandwidth are applied to traffic flows. As a general rule, you must not trust the packet marking from RToWLAN endpoints unless the enterprise has complete control over the endpoints and the applications that are running on those devices. Besides re-marking packets for untrusted devices or applications, administrators can also enable network-based policing and rate limiting to ensure that untrusted devices or applications do not consume too much network bandwidth.

In some deployments, RToWLAN endpoints securely connect to the enterprise network to utilize RToWLAN applications and services. Because these connections traverse the Internet, there is no end-to-end QoS on the IP path. All packet marking is ignored and all traffic is treated as best effort. RToWLAN application performance cannot be guaranteed over these types of connections.

Security

When you deploy wireless endpoints, consider the security mechanisms that are used to control access to the network and to protect the network traffic. Wireless LAN infrastructure and RToWLAN endpoints support a wide range of authentication and encryption protocols, including WPA, WPA2, EAP-FAST, and PEAP. Generally, the authentication and encryption method that you choose for securing the wireless LAN should align with the IT security policies that are supported by both the WLAN infrastructure and the RToWLAN endpoint devices that you deploy.

An authentication and encryption method that supports fast rekeying such as Proactive Key Caching (PKC) or Cisco Centralized Key Management (CCKM) is important for real-time traffic solution deployments. It is critical because it ensures that active voice and video calls and other RToWLAN applications can maintain connectivity and operations as the RToWLAN endpoint is roaming from one access point in the network to another.

Another important security consideration is seamless attachment to the WLAN network. The endpoints must automatically attach to the WLAN network without user intervention to maximize the utilization of RToWLAN applications and services. Certificate-based identity and authentication facilitates an excellent user experience by eliminating user intervention (after initial provisioning) for network connection and minimizing authentication delay. However, deployments where enterprise security policy requires two-factor authentication or one-time passwords, user intervention is required for network attachment. In such cases, access to RToWLAN applications and services gets delayed.

Remote secure attachment

With appropriate security infrastructure and configuration in place, 802.11 RToWLAN endpoints are able to connect to the enterprise from remote locations using public or private 802.11 WLAN networks or Wi-Fi hot spots. While this securely enables RToWLAN application and service delivery for remote attached endpoints, you must consider whether to deliver RToWLAN applications and services over these types of remote connections.

The two key reasons that makes it problematic to enable RToWLAN applications and services over remote secure connections are as follows:

- Nonenterprise 802.11 WLAN

  Public and private 802.11 WLAN networks like wireless hot spots that are found at coffee shops and airports are typically not optimized for real-time traffic applications and do not deliver enterprise-class security or performance. Acceptable RToWLAN solution performance (voice and video quality, connection reliability, and so on) can never be guaranteed over nonenterprise class 802.11 WLANs.

- Internet traversal

  Because remote connections result in real-time traffic traversing the Internet between the enterprise and the endpoint, RToWLAN application performance, including voice and video quality, may be poor. Connectivity across the Internet is never guaranteed and always best effort. There is no end-to-end QoS on the IP path and all packet marking is ignored. All traffic is treated as best-effort. Acceptable RToWLAN solution performance can never be guaranteed over these Internet-based network connections.

Related Topics

Real-Time Traffic over WLAN Radio Frequency Design, on page 21
Real-Time Traffic over WLAN Quality of Service, on page 47
Real-time Traffic over WLAN Security, on page 79
Real-time Traffic over WLAN Roaming, on page 91

RToWLAN solution high availability

High availability is another important factor to consider while planning and deploying RToWLAN. Along with stringent network and radio frequency design requirements for successful deployments of real-time traffic applications and services, you must also consider redundancy and failover for WLAN infrastructure and real-time applications, services, and endpoints.

You must implement the wired network infrastructure in a redundant fashion so that the IP path from the network edge through the data centers and to all locations within the enterprise network infrastructure is maintained even in cases of hardware or service failure. With redundant physical network connections and appropriate network routing and switching configuration, hardwired RToWLAN components, such as call control platforms and other applications servers, are still able to communicate on the network even when components or portions of the network becomes unavailable.

The WLAN infrastructure must be deployed in a resilient manner to ensure continued network connectivity for endpoints even in scenarios where an access point, wireless LAN controller, or authentication system fails. By providing highly available network connectivity, real-time applications and services can continue to function despite isolated infrastructure outages.

Similarly, real-time traffic applications and services as well as the endpoints that service them must be highly available. In the case of real-time voice and video services as provided by Cisco call control, it is imperative that given the failure of a primary call control platform that other platforms or components can continue to

provide these services to endpoints and their users. In addition to network services redundancy, the endpoints and the real-time traffic applications that run on them must be able to automatically fail over to backup service nodes and remain operational.

High availability considerations for RToWLAN solutions deployments, in some cases, dictates the physical characteristics of the enterprise network or deployment.

Single-site or campus RToWLAN deployment

In a single-site or campus deployment, the RToWLAN solution is implemented and operated within a single location or a group of locations within close proximity. The primary consideration for this deployment is to provide resiliency to network connectivity and RToWLAN services like voice and video calling.

Figure 3: RToWLAN single-site deployment, on page 15 shows a single-site or campus deployment that relies on duplication of key components and services to ensure high availability across the RToWLAN solution.

Figure 3: RToWLAN single-site deployment

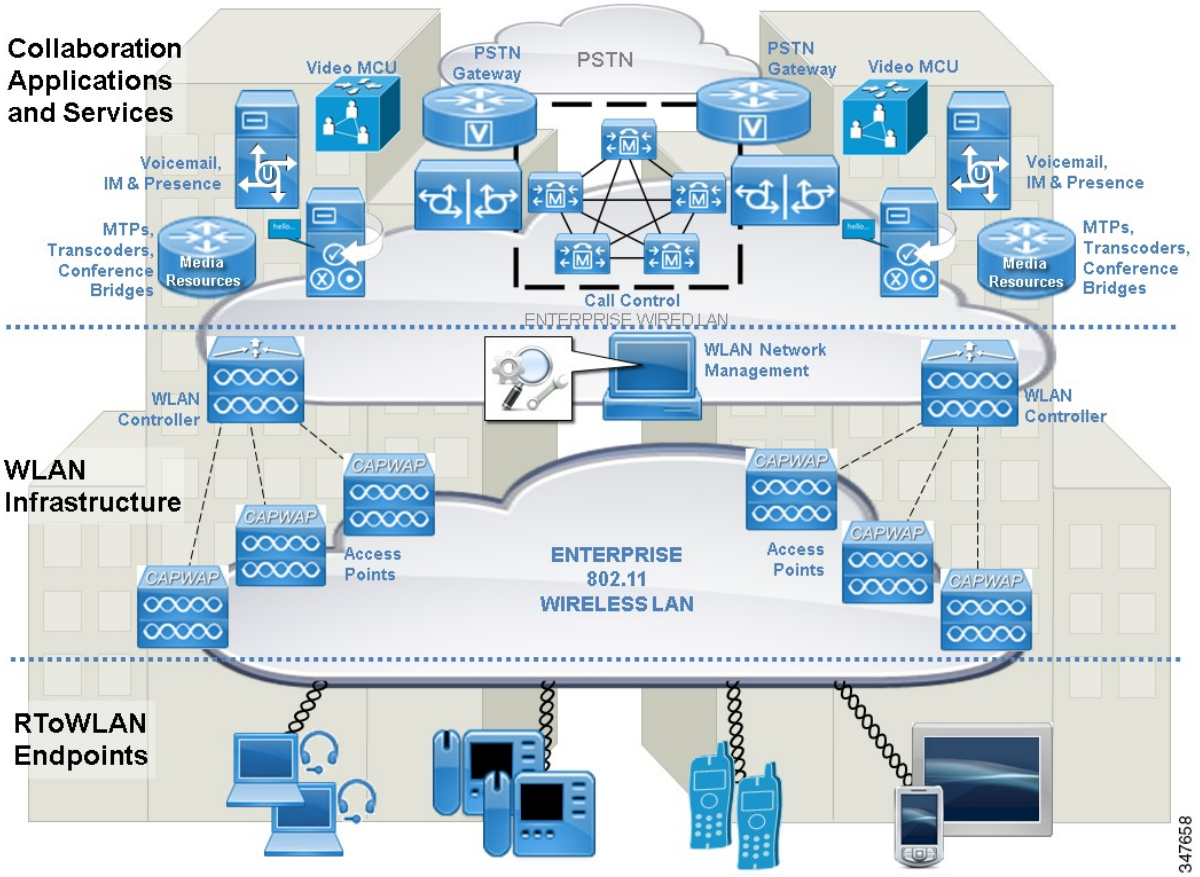

Note The preceding figure does not show the wired network infrastructure redundancy but you must assume that it exists.

It also shows a split data center between two buildings in a campus location with duplication of all key components and services, which are as follows:

- Wireless network components: Multiple WLAN controllers and access points.

  Deploying redundant wireless network infrastructure components ensures that wireless network connectivity for RToWLAN endpoints is highly available. It also ensures that the devices can continually have access to network-based RToWLAN applications and services both when stationary and when moving or roaming within a location.

- Collaboration components: Multiple call control platforms or nodes, media resources (conference bridges, media control units, and so on), PSTN gateways or border controllers, and application servers.

  Deploying redundant call control platforms, media resources, PSTN connections, and application servers ensures that RToWLAN applications and services are highly available and that RToWLAN endpoints can continually use these applications and services.

Distributed RToWLAN deployment

In a distributed deployment, the RToWLAN solution is implemented and operated at multiple sites or locations, and RToWLAN devices are distributed throughout the network. In these type of deployments, providing resiliency to network connectivity and RToWLAN services like voice and video calling is still the primary consideration. However, there are specific considerations apply for distributed deployments, related to site interconnectivity.

Figure 4: RToWLAN Distributed Deployment, on page 17shows that with distributed deployments, key WLAN and collaboration components and services must be replicated locally at each site to provide persistent network connectivity and persistent access to RToWLAN applications and services across the enterprise.

Figure 4: RToWLAN Distributed Deployment

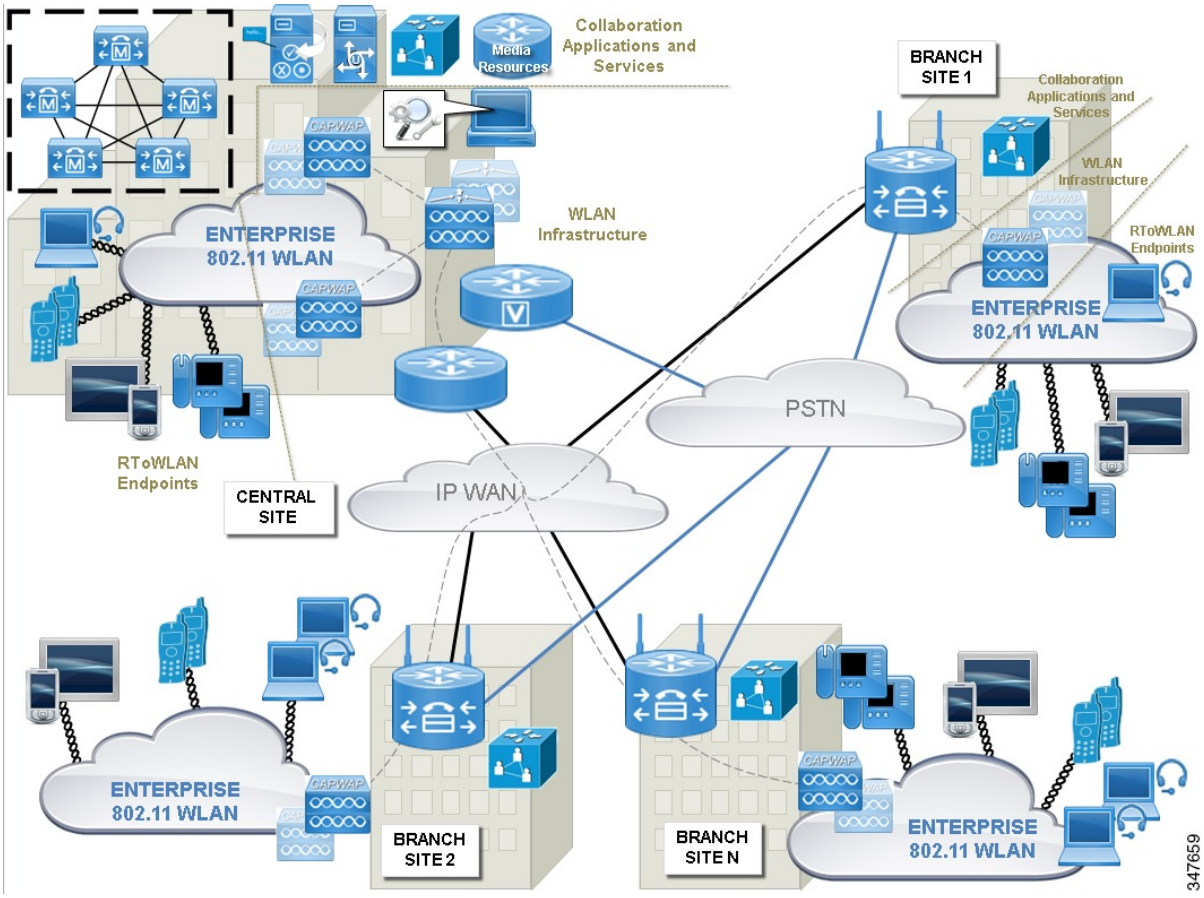

Access points (APs) are deployed redundantly throughout the enterprise at each site based on density and site survey requirements. In the case of wireless LAN controllers (WLCs), you have two high availability deployment options:

- Central site WLCs control and manage APs at all locations providing centralized management and control.

  In the case of IP network failures between a branch and the central site, the local APs at the branch continue to service local RToWLAN endpoint devices, providing network connectivity and authentication services with cached credentials from the central site. This type of high availability scheme works best for smaller branch sites with a limited number of APs and wireless devices.

- Local WLCs deployed at each branch site control and manage APs at the local site.

  In the case of IP network failures between a branch and the central site, all WLAN network services continue to be provided locally. It is also possible to leverage centralized authentication services in this type of deployment, with the local site WLC caching credentials from the central site to provide local authentication services during IP network outages. This type of high availability is better suited for larger

branch sites with many APs and wireless devices. Even in the case of distributed or local WLCs, in most deployments, the overall WLAN network management component and application typically remain centralized as shown in Figure 4: RToWLAN Distributed Deployment, on page 17.

Note In deployments where branch sites vary in size or the required service level at each location differs, it is possible to have hybrid deployments, where some locations rely on centralized management and control of wireless APs and other locations rely on distributed or local control.

Figure 4: RToWLAN Distributed Deployment, on page 17 depicts both WLC high availability deployment options. The gray dashed lines between the central site WLC and all APs including the branch site APs represent centralized WLC control. Each branch site location router is capable of providing local WLC capabilities, as indicated by the wireless antennas.

Collaboration call control platforms and other components must also be replicated at each enterprise site to provide persistent access to RToWLAN applications and services. Figure 4: RToWLAN Distributed Deployment, on page 17 shows the replication of critical collaboration components including call control, PSTN connectivity, and media resources as represented by the voice enabled router and the media control unit (MCU) at each branch site. As with WLC high availability for distributed deployments, collaboration call control can also be centralized within the central site or distributed at all locations. Figure 4: RToWLAN Distributed Deployment, on page 17 depicts centralized call control with distributed backup call control provided at each branch site in the case of IP network failures between the branch and central site.

Another important consideration for RToWLAN deployments and distributed multisite deployments is RToWLAN device mobility. In a multisite deployment, movement between locations is common, because of the mobile nature of most RToWLAN endpoints. The collaboration call-control platform should dynamically track the location of RToWLAN endpoints as they move between enterprise sites. Based on the IP address of the device or other identifying information, the call control application should determine the devices location and adjust call routing, PSTN egress point, and codec and media resource selection as appropriate.

Another important consideration with distributed RToWLAN deployments is call admission control. Call admission control is a collaboration call control feature that ensures that bandwidth between enterprise sites is not oversubscribed by voice or video call traffic. Oversubscription of bandwidth on the connections between sites leads to poor voice and video quality, delayed call set up, and even dropped calls. Given the limited bandwidth and throughput that is available on links between enterprise sites which are usually lower speed, call admission control ensures that sufficient bandwidth is available on the IP path to setup and maintain a quality voice or video call. If sufficient bandwidth is not available, the call control system denies the call setup or reroutes the call using the local site PSTN connection. While call admission control is not unique to RToWLAN, it is an important consideration in distributed multisite collaboration deployments.

RToWLAN solution capacity planning

RToWLAN deployment scalability is a major design consideration when you implement real-time traffic applications and services. Failure to provide sufficient network and call processing capacity may result in service or functionality outages that prevent endpoints from associating, authenticating, registering, making or receiving voice and video calls, or leveraging other collaboration applications.

The WLAN infrastructure must provide sufficient client attachment as well as bandwidth capacity to ensure endpoints are able to actively connect to the WLAN and successfully make high-quality voice and video calls. In particular, the number of simultaneous voice or video bidirectional traffic streams per WLAN channel cell

is a critical capacity consideration, which, along with the number of associated endpoints, ultimately determines deployment device densities and potential user call rates.

Apart from WLAN infrastructure endpoint and bandwidth capacity, you must also consider the capacity of the collaboration system, which enables real-time traffic applications and services. With Cisco call control, each call control platform or node has a finite endpoint and call volume capacity. Likewise, media resource platforms and MCUs have finite call or session capacity. You must deploy sufficient endpoint and call volume capacity by adding the appropriate number of call control application nodes to provide services for the required number RToWLAN endpoint users.

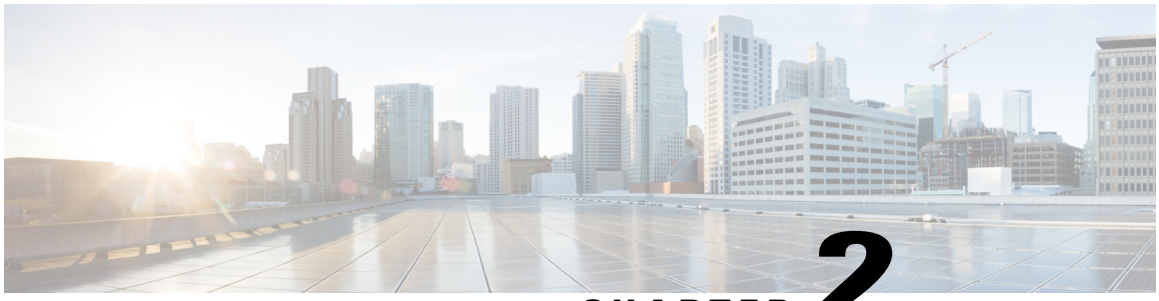

CHAPTER 2

Real-Time Traffic over WLAN Radio Frequency Design

The chapter describes, in general terms, the radio frequency (RF) plan and design considerations for RToWLAN deployment. The other factors that affect the RF plan and design considerations are endpoint capabilities, local conditions, and regulations. This chapter presents typical deployment scenarios to illustrate RF-related processes and considerations.

High availability

High availability (HA) is an important factor when you consider a system plan that includes RToWLAN. For an RToWLAN deployment, you can apply the same HA strategies that you use in wired networks, to the wired components of the RToWLAN solution. A unique factor to RToWLAN availability that you must consider is the RF coverage HA, that is, providing RF coverage that is not dependent on a single WLAN radio. The primary mechanism to provide RF HA is *cell boundary overlap*.

An overlap of 20 percent means that 80 percent of a given access point (AP) cell is also covered by other APs at the recommended signal levels, while in the other 20 percent of the cell RToWLAN calls are of degraded quality, but are still available. The Cisco Unified Wireless Network Coverage Hole algorithm amplifies the RF HA coverage, which detects if WLAN clients are experiencing poor signal-to-noise ratio (SNR) values and causes the power of APs to increase as needed to rectify SNR issues.

Note In deployments planning to rely on the Coverage Hole algorithm, the planning must consider whether an AP is going to increase its power level to adjust for a hole. Therefore, the maximum power of the RToWLAN endpoint, which can be lower than the maximum power of the AP, needs to be considered when configuring the initial AP power level.

Related Topics

Coverage hole algorithm, on page 23

Capacity planning

Capacity planning is another important parameter in an RToWLAN plan. Call capacity is the number of simultaneous RToWLAN calls that an area can support. The number of calls can vary depending on the RF environment, the RToWLAN endpoint features, and the WLAN system features.

Cisco strongly recommends the use of the 5 GHz spectrum channels for RToWLAN design, especially when video is part of the media that is to be transmitted over the WLAN.

The following table provides three examples of the best case scenario approximate maximum capacity per access point or channel for video calling when endpoints use a WLAN that provides optimized WLAN services (such as the Cisco Unified Wireless Network). These examples consider an RToWLAN that uses 5 GHz channel of the 802.11n WLAN standard, with no Bluetooth, and with no channel bonding.

Table 3: Video over WLAN Call Capacity

Estimated maximum number of simultaneous bidirectional Video Calls	Resolution/Bitrate	WLAN Standard	Data Rate/MCS Index
7	H.264 720p/2500 Kbps	802.11n	MCS 7 (40 MHz Channels)
4	H.264 720p/2500 Kbps	802.11n	MCS 4 (40 MHz Channels)
1	H.264 720p/2500 Kbps	802.11n	MCS 1 (40 MHz Channels)
16	H.264 360p/400 Kbps	802.11n	MCS 7 (40 MHz Channels)
12	H.264 360p/400 Kbps	802.11n	MCS 4 (40 MHz Channels)
8	H.264 360p/400 Kbps	802.11n	MCS 1 (40 MHz Channels)

Because the 5 GHz spectrum features less noise and interference, there can be greater capacity with the higher carrier frequency implementation. The additional nonoverlapping channels that are available in the 5 GHz spectrum also provide higher call capacity for a given area. In addition, it is recommended to use the 40 MHz channels for video because they provide increased capacity. Consider as well that these maximum approximate numbers vary according to your particular environment noise, coverage, attenuation, Bluetooth utilization, channel utilization, how many spatial streams the endpoint in question supports, and the mix of clients in the cell.

The following table provides two examples of the estimated maximum capacity per access point or channel for voice calling when endpoints use a WLAN that provides optimized WLAN services. These examples consider an RToWLAN that uses a 5 GHz channel of the 802.11a WLAN standard and with no channel bonding.

Table 4: Voice over WLAN Call Capacity

Estimated maximum of simultaneous bidirectional Voice Calls	Audio Codec/Bit Rate	WLAN Standard	Data Rate
20	G.711-G.722/64 Kbps	802.11a	12 Mbps
27	G.711-G.722/64 Kbps	802.11a	24 Mbps or higher

Note The call capacities mentioned above are per nonoverlapping channel because, here, channel capacity becomes the limiting factor and not the number of APs. These maximum call capacity numbers are provided for general planning purposes. You must use the call capacity that is specified by the actual RToWLAN endpoint for deployment, because this is the supported capacity of that endpoint. Additionally, Cisco recommends that you use 40 MHz channel when you utilize the 5 GHz band. For more information to determine accurate values for call capacity planning, see the endpoint documentation.

Coverage hole algorithm

In deployments that plan to rely on coverage hole algorithm, during planning, you must consider whether an AP is going to increase its power level to adjust for a hole. Therefore, the maximum power of the RToWLAN endpoint, which can be lower than the maximum power of the AP, must be considered when you configure the initial AP power level. For example, if the RToWLAN endpoint has a maximum power level of 16 dBm and the initial AP power level is 16 dBm, increasing the AP power to cover an RF hole does not help a RToWLAN client in that hole, because the RToWLAN will not be able to increase its power level any further to compensate for the increased power from the AP. In this example, for the hole coverage to be effective, the initial AP power level should be 13 dBm or less such that both the AP and the RToWLAN have sufficient room to increase power to cover an RF hole.

If you place APs with a design goal to support media-rich and real-time applications, the population of APs relatively changes from dense to highly dense. The remote management module (RMM) power level assignments of levels 3 and 4 are of medium density and the power levels of 5 through 8 are highly dense. If the AP placement is such that the RMM has assigned power levels of 1 or 2, the coverage is not considered dense. At AP power level of 1, the APs surrounding an AP which drops out of the network are not able to

increase their power. They will not be able to cover the hole that was created by an AP that is going off line. If the power levels of the surrounding APs was 2, the surrounding APs would increase the power level to 1.

If minimum density is a design goal, then you must do the survey process and initial coverage evaluations with the actual APs and the actual antennas that are to be installed. Then set the maximum transmit power in dBm in RMM to the maximum transmit power of the weakest client. The initial AP placements should be such that their coverage area is no greater than the coverage area of the weakest client. Set the dBm value in RRM to equal the dBM transmit power of the weakest mobile client so that the coverage of the APs match the coverage of the clients.

The following figure shows the transmit power control settings of a Cisco WLAN.

Figure 5: Transmit Power Control for Coverage Hole Algorithm

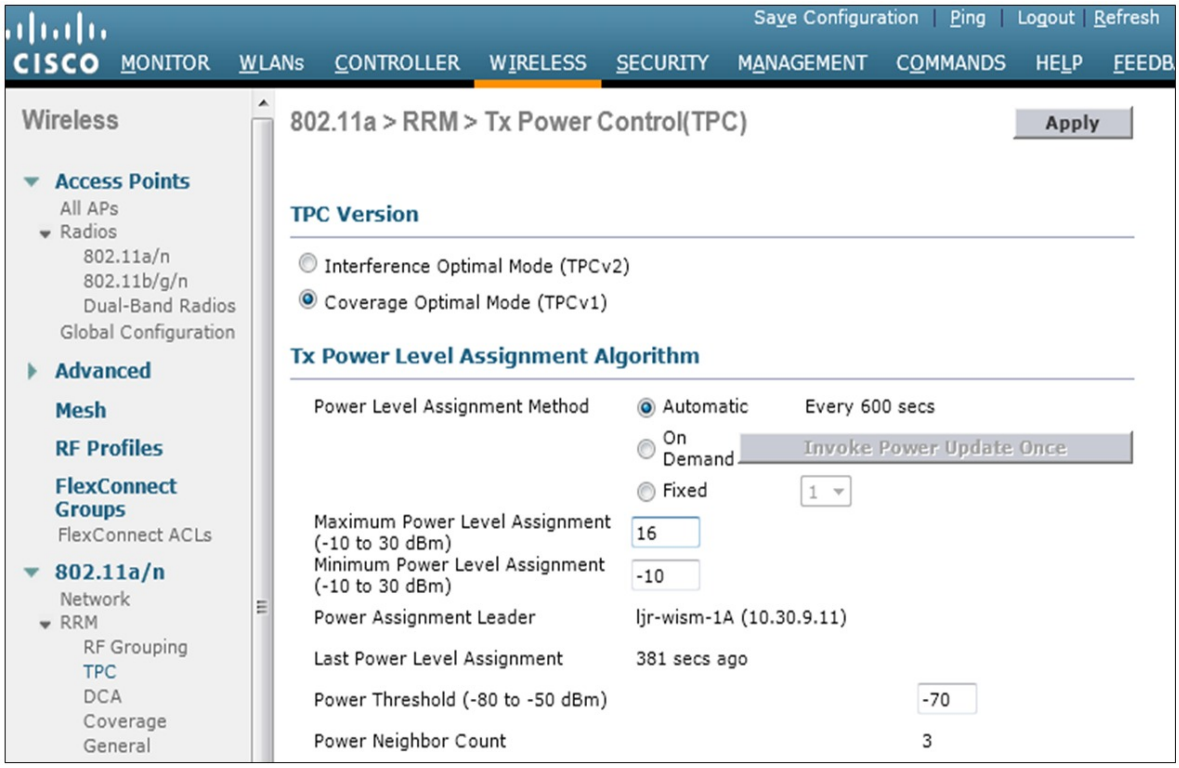

All clients that use 802.11g, 802.11a, 802.11n, or 802.11ac take advantage of client link downstream and maximum ratio combining (MRC) upstream. This is a dynamic Wi-Fi per client quality of signal enhancement and can be effective when there are coverage holes from AP outages.

Design considerations

This section describes important elements that you must consider when you design RF coverage for wireless endpoints using RToWLAN.

General AP guidelines for rich media

The packet loss and jitter requirements of rich media and the increased mobility of RToWLAN endpoint users place demands on connection quality and coverage that are beyond that of a typical WLAN data deployment. Although later generations of WLAN equipment and software can provide further RToWLAN improvements, the foundation of a successful RToWLAN deployment depends on radio frequency planning, designing , and implementing. It is important that you design, plan, implement, operate, and maintain the WLAN RF environment to make the RToWLAN deployment successful. The processes, guides, heuristics, and tools that are used for a WLAN data deployment do not help to deliver a successful RToWLAN deployment.

The general RToWLAN guidelines for an optimal RToWLAN network are as follows:

- Cell overlap of at least 20 percent and overlap of approximately 20-30 percent for critical communications for industries like healthcare, where a WLAN Data design can use an AP cell overlap of 5-10 percent.
- Cell boundary of -67 dBm.

Note Cisco strongly recommends the use of 5 GHz spectrum channels for RToWLAN design, especially when video is part of the media that is to be transmitted over the WLAN. If possible, use the 40 MHz channels instead of the 20MHz.

- If you use 802.11n AP platforms, use ClientLink/beamforming to optimize the WLAN performance.
- For optimal performance of voice or video deployments, data rates that are below 12 Mbps should be disabled (including MCS 0).

Note The RF characteristics of RToWLAN endpoints vary, and can affect the WLAN design and capacity greatly. If you are planning a deployment of an RToWLAN endpoint with RF deployment requirements that do not match with those that are presented in this chapter, then you must follow the endpoint guidelines. Although endpoint recommendations vary, the general principles and issues that are discussed in this chapter still apply with changes in cell sizes.

2.4 GHz network design

Cisco recommends that you use the 5 GHz spectrum channels for RToWLAN design.

The IEEE 802.11b/g channel set defines a total of 14 channels. Each channel is 22 MHz wide, but the channel separation is only 5 MHz. This leads to channel overlap such that signals from neighboring channels can interfere with each other.

In a 14-channel DS system (11 usable channels in the U.S.), there are only three nonoverlapping (and thus, noninterfering) channels, namely 1, 6, and 11, each with 25 MHz of separation. This channel spacing administers the use and allocation of channels in a multi-AP environment, for example an office or campus. APs are

usually deployed in a cellular fashion within an enterprise, where adjacent APs are allocated nonoverlapping channels. See the following figure, which illustrates 2.4 GHz channel allocations.

Figure 6: 2.4GHz Channel Allocations

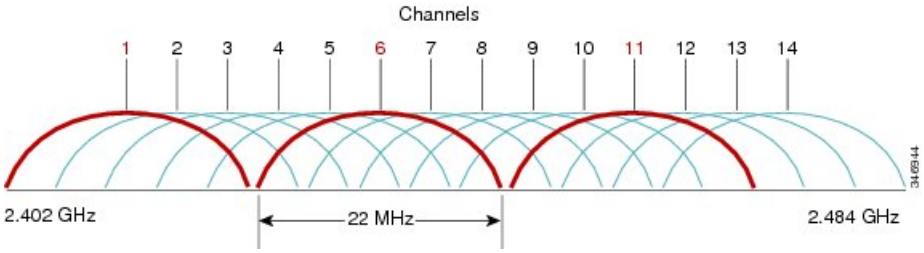

IEEE 802.11b provides data rates of 1, 2, 5.5, and 11 Mbps. IEEE 802.11g provides data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps in the 2.4 GHz band, in the same spectrum as IEEE 802.11b. IEEE 802.11g is backward-compatible with IEEE 802.11b with a single AP providing WLAN access for both IEEE 802.11b and IEEE 802.11g clients.

Co-channel interference considerations

As mentioned in the preceding section, there are only three nonoverlapping channels in U.S. 2.4 GHz spectrum. Therefore, it is difficult when you try to deploy APs and ensure that APs on the same channel do not receive signal from other APs on the same channel. The AP coverage radius changes with the client bit rates that are supported, and the boundary that is created by this radius is considered the AP boundary.

In reality, co-channel interference becomes complicated because the AP influences the WLAN RF environment around it for a much larger distance than just the bit rate boundary. This is because the RF energy from the AP, although too low to be demodulated in to a WLAN frame, is strong enough to cause an IEEE 802.11 radio to defer sending. In addition to the AP influence of the RF environment, the clients that are associated with that AP extend the range of the RF energy that is associated with that APs cell even further.

The IEEE 802.11 MAC is a Carrier Sense Multiple Access-Collision Avoidance (CSMA-CA) algorithm, and the Carrier Sense performs a Clear Channel Assessment (CCA) before it attempts to send an IEEE 802.11 frame. The CCA mechanism is specified for each IEEE 802.11 physical layer; it is triggered either by a simple raw energy level, and Physical Layer Convergence Protocol (PLCP) header power levels, or carrier detection. The CCA of an IEEE 802.11 radio does not vary with the bit rates that are being used and is not, generally, user-configurable.

The impact of CCA deferrals on an AP WLAN from IEEE 802.11 radios that are not part of that AP WLAN is called co-channel interference. As co-channel interference results in delays in sending frames, it causes increased jitter and delay during RToWLAN calls. Although WLAN QoS prioritizes WLAN traffic, this occurs after the CCA and therefore prioritization does not overcome the jitter and delay that are introduced by CCA.

An RToWLAN endpoint must have a power level boundary of -67 dBm and a separation between adjacent AP channels of -86 dBm. The -67 dBm requirement is to minimize packet loss, and the -86 dBm requirement provides separation between adjacent channel cells to minimize co-channel interference from other AP cells on the same channel.

The following figure shows an example of the two boundaries that are created by the -67 dBm and -86 dBm requirements, based on standard RF loss formulas for an open office environment.

Figure 7: Bit Rate and Co-channel Interference Boundaries of an AP

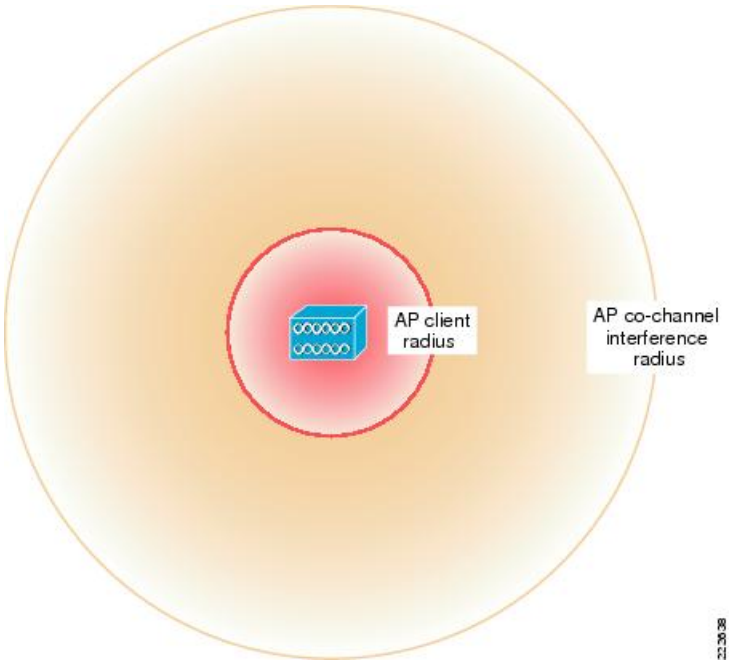

This RF environment, which would give an AP a client radius of 43 feet, which gives an AP co-channel interference radius of 150 feet using standard antenna gain (2 dB), and an AP output power of 16 dBm (40 mW). Different RF environments, AP powers, and antennas provide different client and co-channel interference radii, but the principles that are discussed in this chapter still apply.

Note The output power that you choose for the AP must align with the RToWLAN endpoint capabilities and deployment requirements. For example, the 9971 collaboration endpoint has a maximum output power of 40 mW (16 dBm) when using 802.11a. An AP power greater than 40 mW must not be used for a 9971 collaboration endpoint deployment if 802.11a is used. In circumstances where the Cisco Unified Wireless Network Hole Coverage mechanism that is expected to provide RToWLAN coverage in the event of an AP outage, an AP power of less than 40 mW (using the 9971 collaboration endpoint as an example) must be used for AP planning to allow the APs covering an RF hole to operate in a range that is suitable for the RToWLAN endpoint.

One additional advantage of using a lower AP transmit power is a proportional decrease in the co-channel interference radii. In the preceding example, a 40 mW (16 dBm) transmit power gives a co-channel radius of the 150 feet and a client radius of 43 feet. A decrease in the power to 20 mW (13 dBm) reduces the co-channel radius to 130 feet and the client radius to 38 feet, and also reduces the co-channel interference proportional to the co-channel interference that is generated by an AP.

The RF co-channel interference radius of an AP as well as a WLAN client contribute to the co-channel interference. See the following figure.

Figure 8: Single Client Co-channel Interference Radius

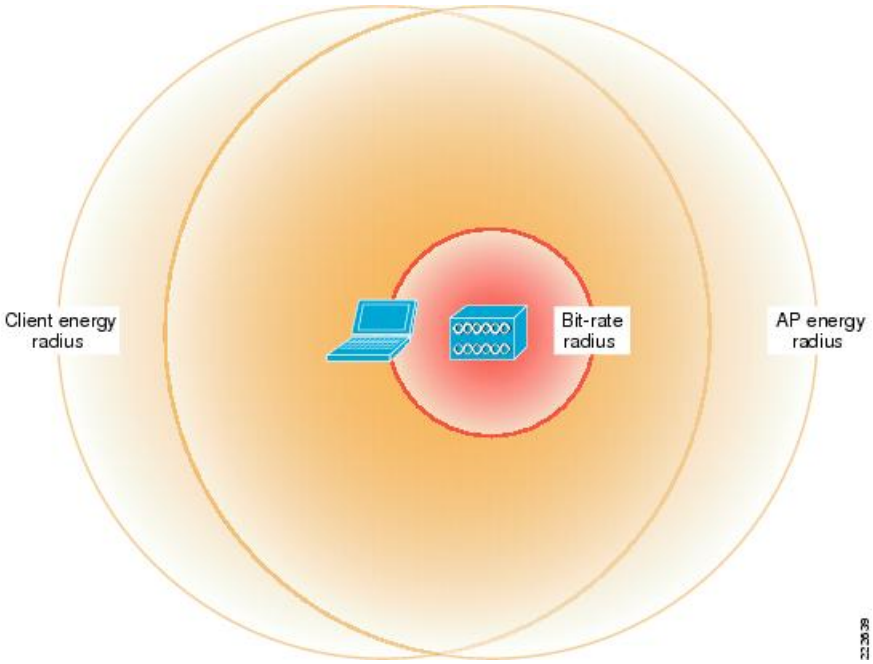

The client co-channel interference is better illustrated in the following figure provided a client or clients can be located anywhere on the bit rate radius perimeter. With the 43 foot bit rate radius and 150 feet AP co-channel interference radius of the previous calculations, 193 foot is the new client co-channel interference radius.

Figure 9: Complete Client Co-channel Interference Radius

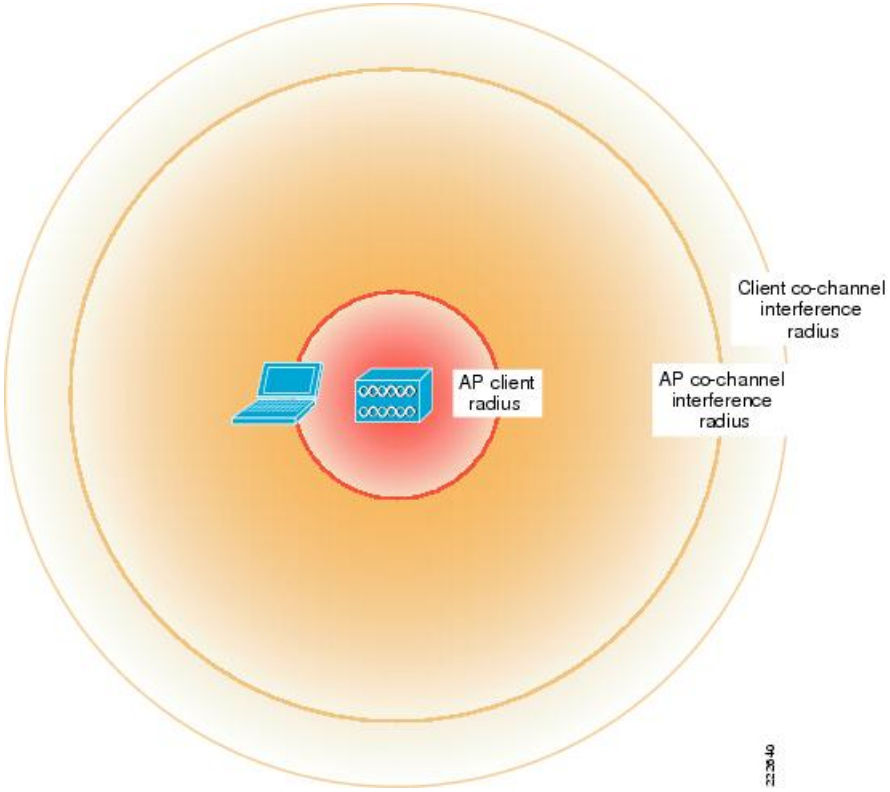

Note The 193-foot radius represents close to a worst case because the WLAN client is not normally in an equivalent location to an AP and will likely suffer greater signal attenuation due to obstacles.

Bit rate impact on co-channel interference

The AP client radius in the example results in a nominal bit rate for the RToWLAN endpoints of approximately 24 Mbps or greater, depending on noise. You can extend the AP client radius further by lowering the bit rates. But, this is not recommended for the following reasons:

- Lowering the bit rate extends the AP client radius, but also increases the client co-channel interference radius, increasing the area that has the RToWLAN call capacity of only a single AP. Furthermore, lower bit rate reduces video call quality.
- Lowering the bit rate reduces the overall call cell capacity, because lower bit rate packets consume more time, and transmit fewer packets.

RToWLAN call quality is sensitive to data rate shifting. The decision to make a data rate shift is normally the result of being unable to send frames at the data rate that was previously used, which is determined by

sending a frame multiple times without receiving an acknowledgment for that frame. This increases the delay and jitter during an RToWLAN call.

Some clients can utilize the traffic stream rate set (TSRS) IE to use a subset of enabled data rates (for example, 12-24) to help first transmission success.

20 percent cell overlap

It is recommended that you use an AP cell overlap of at least 20 percent for RToWLAN deployments. This ensures RToWLAN endpoint can detect and connect to alternative APs, when it is close to the cell boundary. It also allows an RToWLAN client to change AP associations with a minimum interruption to a call, by minimizing the amount of data rate shifting and retransmission at a cell boundary for a given RToWLAN client.

The 20 percent overlap requirement means that APs are spaced closer together than the two-times-70 feet distance suggested by the cell boundary. The area of overlap between two circles of radius is equal to 1. d is the distance between the centers of each circle. For an area of 20 percent, the value of d is 1.374 for a standard radius of 1, or 59 feet between APs for our 67 dBm boundary.

Note The other common distance values that are used are 10 percent (1.611), 15 percent (1.486), 25 percent (1.269), and 30 percent (1.198).

The following figure illustrates 20 percent AP overlap.

Figure 10: APs with 20 Percent Overlap

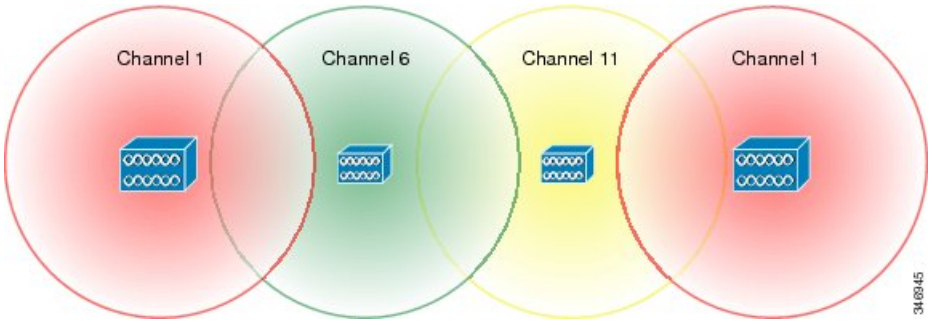

Co-channel interference and 20 percent AP cell overlap

The following figure shows an APs with 20 percent overlap and their co-channel interference boundaries.

Figure 11: AP with 20 Percent Overlap and Co-channel Interference Boundaries

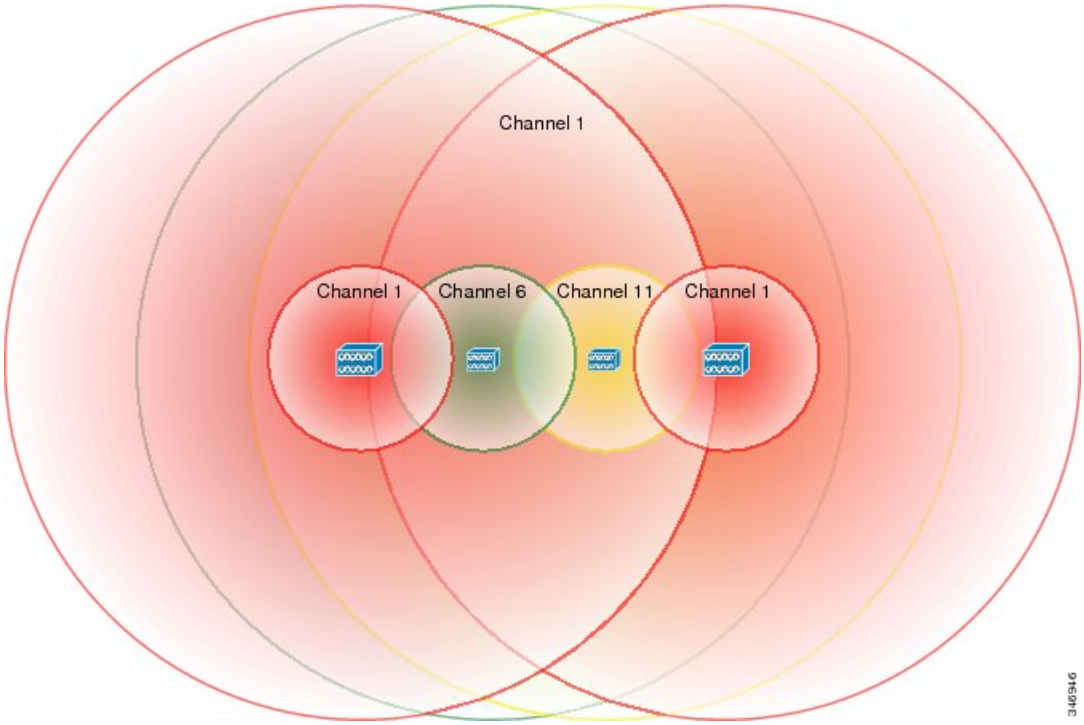

The co-channel interference boundary for one of the APs using channel 1 overlaps with an AP using the same channel. In this case, co-channel interference occurs in an RToWLAN deployment. You must also note that the combined effect of the 20 percent overlap requirement for reliable roaming between AP cells and the impact of co-channel interference is a reduced per RToWLAN channel cell call capacity over a given area.

Note It is not effective to reduce overlap in order to reduce co-channel interference. This results in poor roaming performance and reduces user satisfaction. In contrast, you can address call capacity while you plan and design.

Existing WLAN data deployments (that initially used lower-power cell-boundaries and less overlap) that are changed to match recommended power boundaries and overlap for RToWLAN can experience application issues for time-sensitive applications. It is difficult to predict which applications can be affected by the WLAN changes, because the actual effect depends on the application implementation. In general, custom applications that require keepalive timeouts are most likely to be affected, and must be validated in the new environment to ensure that their timers require no adjustment.

Deployment examples

The AP layout within a building depends on the building construction and shape, and the WLAN coverage requirements in that building. Due to different effects of implementation-specific variables, there is not a single recommended deployment for the number of APs that you must deploy nor a single solution to determine

the effect of co-channel interference. The following sections describe the design process with examples that illustrate deployment types:

- Single-floor building deployment example
- Multifloor building deployment example

Single-floor building deployment example

The illustrations for the single-floor deployment with 20 APs and 15 APs do not show the location of building exits. It is critical that there is coverage around the building exits if not between the buildings.

The following figure shows a rectangular building with a dimension of 285 feet x 185 feet that will require 20 APs to give complete coverage.

Figure 12: 2.4 GHz Single Floor Deployment with 20 APs

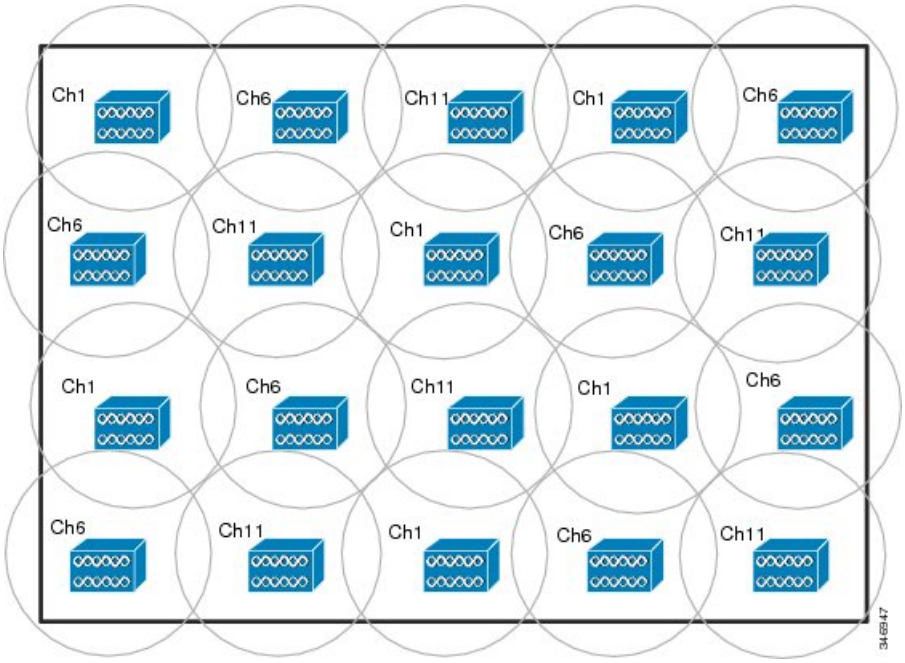

A WLAN data deployment with the same AP boundary and plan will be able to use only 15 APs, but this gives small coverage gaps and less overlap, as shown in the following figure. One of the characteristics of

RToWLAN deployments is that users are more mobile and find coverage gaps that WLAN data clients cannot find. Therefore, it is preferred that you use a 20 AP deployment.

Figure 13: 2.4 GHz Single Floor Deployment with 15 APs

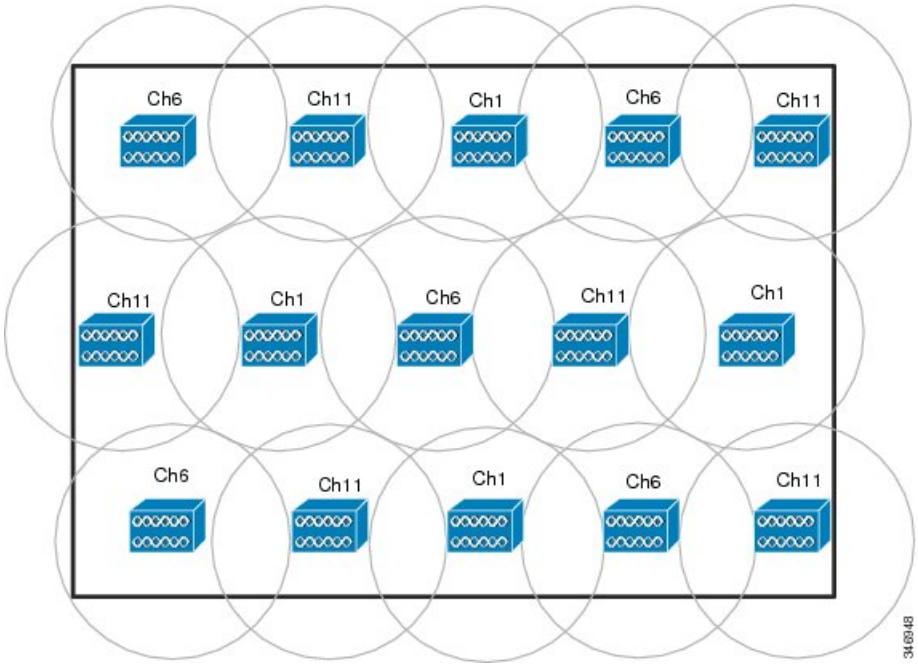

The following figure shows the co-channel interference radius of an example AP where the co-channel interference radius extends for the entire building. This means the APs that are using channel 1 are effectively sharing channel capacity. The six channel 1 APs have increased the coverage over single AP by six times, but not increased the capacity by the same ratio, and might not increase the capacity significantly in comparison with single AP. This is applicable for the APs on other channels. Because of co-channel interference, the call capacity of the floor is equivalent to something above the capacity of three independent APs, but not

approaching the capacity of 20 APs. This is the primary reason to address RToWLAN call capacity in terms of the number of calls per channel, and not the number of calls per AP. Channel capacity is the limiting factor.

Figure 14: 2.4 GHz Single-Floor Co-channel Interference

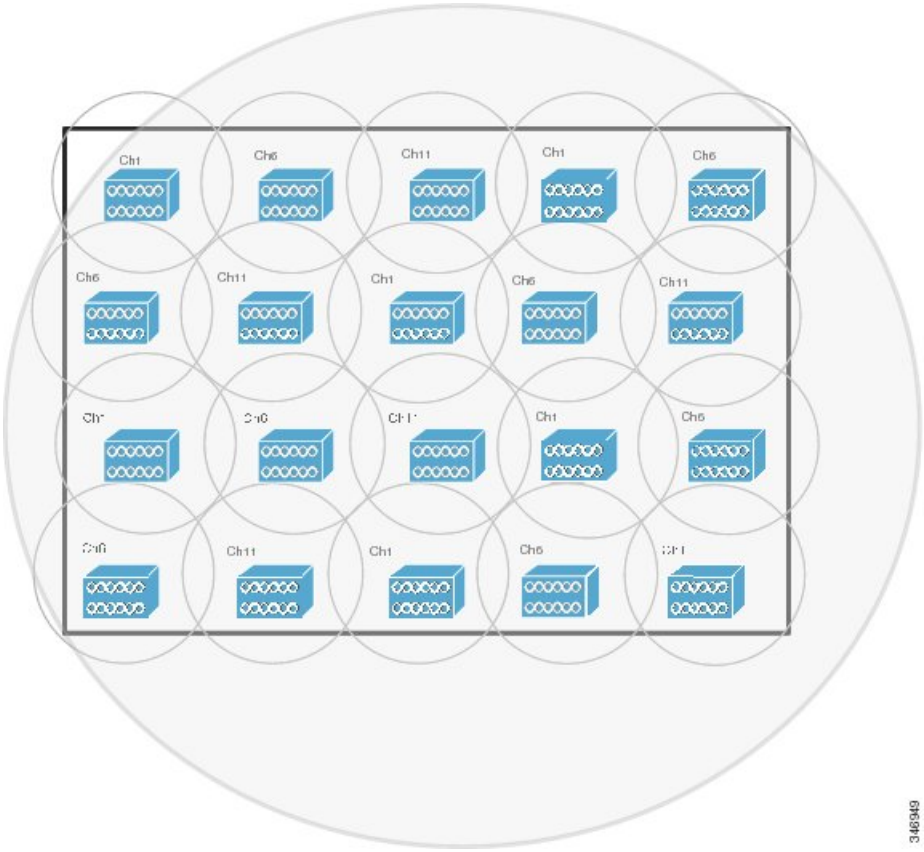

Note Given the security concerns, it is recommended that you do not extend cell radius outside physical building boundaries, except in scenarios where wireless connectivity is required outside the building. For example, wireless coverage is required between buildings in a campus deployment.

Multifloor building deployment example

Note The examples that are provided in the multifloor building deployment only illustrate the concept of multifloor channel assignments and co-channel interference between floors. It is recommended that you do not use these channel implementation examples for real-time deployments.

In a multifloor building, RF energy can travel between floors and, as part of RF planning, the channels are staggered from floor to floor to minimize the co-channel interference between floors, as shown in the following figure. When you consider the co-channel interference radius of an AP, you must note that the signal path between the floors is different from that on the same floor (there is often a piece of reinforced concrete in the

between-floor path). You must take this into account when you consider co-channel interference radius of an AP.

Figure 15: 2.4 GHz Multifloor Channel Assignments

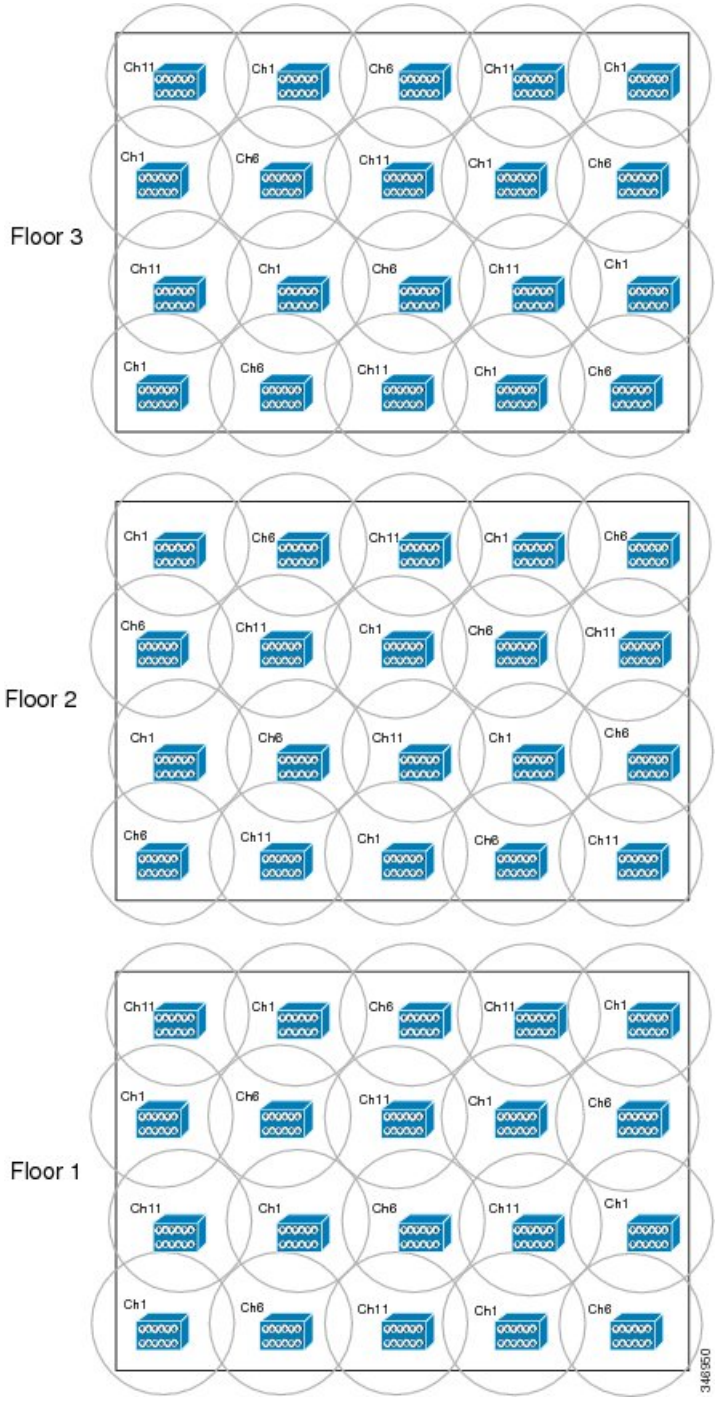

The following figure shows an example of the co-channel interference radius of APs on different floors where floor 2 is the same layout as the single-floor example, and floor 1 and floor 3 show the co-channel interference radius on the floors above and below. In the following figure, the co-channel interference between floors is

still significant and it is reasonable to assume that the capacity across the three floors can be the equivalent of six or seven APs, but is not close to that of the 60 APs that are deployed.

Figure 16: 2.4 GHz Multifloor Building Showing Co-channel Interference

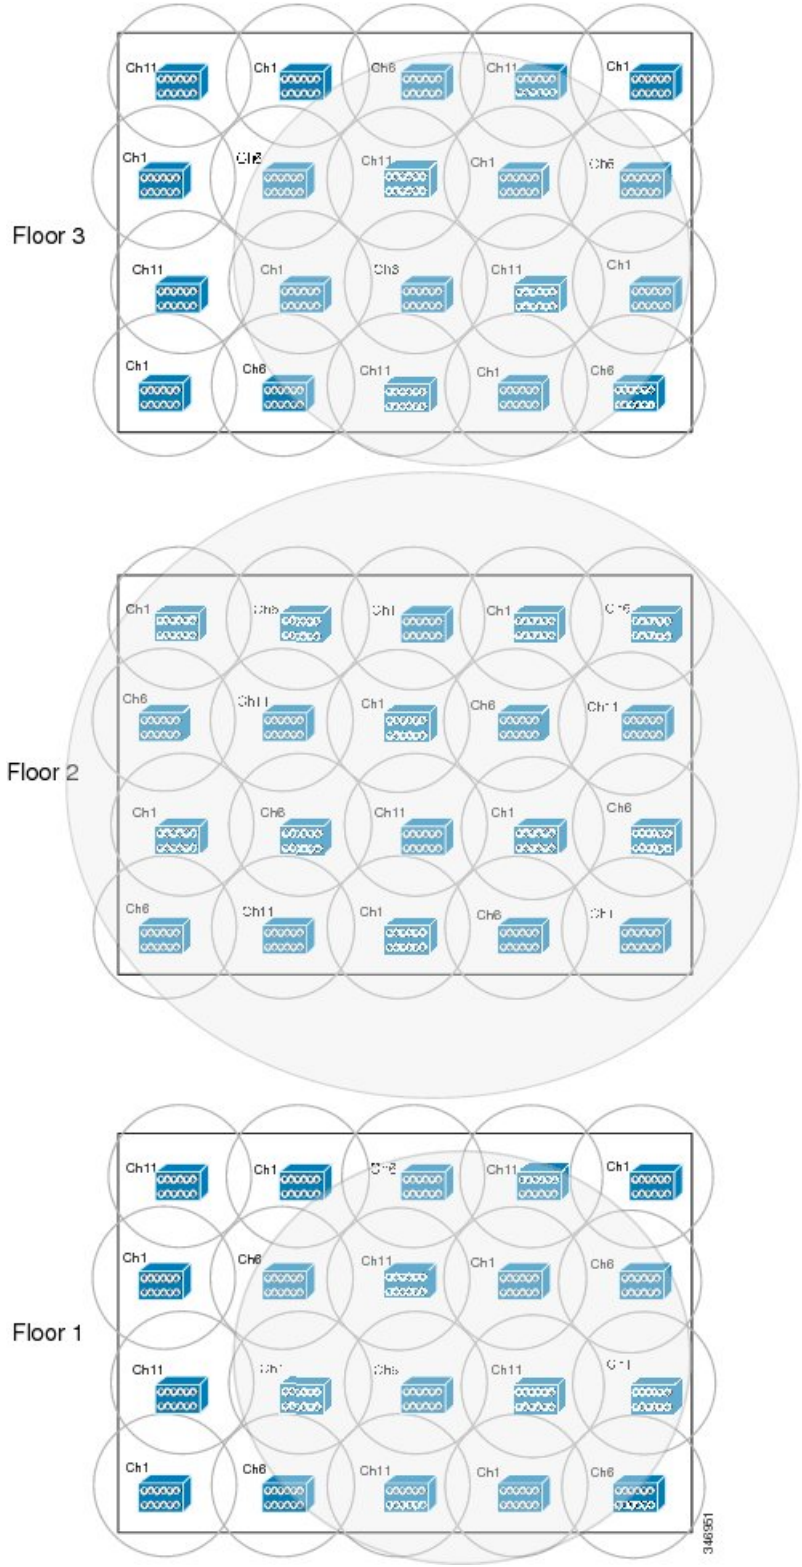
Floor 3
Floor 2
Floor 1
346951

Location-based services design considerations

The signal-level requirements of IEEE 802.11 location-based services are similar to those on RToWLAN, but the AP placement requirements are different. For example, in Figure 12: 2.4 GHz Single Floor Deployment with 20 APs, on page 32, the AP placement satisfies the requirement of location-based service (LBS) deployment. In this example environment, many APs are deployed on the perimeter and at the core of the building. An additional set of APs might be required. The AP placement requirements of LBS can result in addition of more APs, depending on the shape and size of the building. The addition of more APs for LBS introduces an additional level of co-channel interference due to the additional IEEE 802.11 management traffic that is associated with additional AP. However, given the existing co-channel interference, the difference is not significant. The key point about RToWLAN deployment is that the addition of more APs do not contribute to additional capacity in the 2.4GHz band due to co-channel interference.

Auto-RF significance

The co-channel interference in the 2.4 GHz band affects the RToWLAN channel cell call capacity, and therefore you must consider this limitation while you plan, design, and operate an RToWLAN network. All the examples in this chapter assume Auto-RF is enabled, and that an AP changes channels to minimize interference and provide an optimal channel plan. With this assumption, the call capacity of the deployment is the equivalent of three times the capacity of a single AP. The auto-RF cannot do much to address the effects of co-channel interference in the 2.4 GHz channel because the limiting factor in the 2.4 GHz band is the three nonoverlapping channels. Auto-RF tunes AP power levels that can reduce co-channel interference by reducing power levels. But this power level adjustment must be balanced against the signal-level and coverage requirements of the RToWLAN deployment. It is best that you use the 5 GHz band of IEEE 802.11a standard to achieve greater capacity and a higher return on investment of the deployed APs.

5 GHz network design

This section describes the following considerations for implementing a 5 GHz RToWLAN deployment:

- IEEE 802.11a physical layer
- IEEE 802.11a channels
- IEEE 802.11a operating frequencies and data rates
- IEEE 802.11a and RToWLAN deployments

IEEE 802.11a physical layer

The IEEE 802.11a standard defines the requirements for the physical layer (of the OSI model), operating in the 5 GHz unlicensed national information infrastructure (UNII) frequency band, with data rates ranging from 6 Mbps to 54 Mbps. It uses orthogonal frequency division multiplexing (OFDM), which is a multicarrier system (uses 52 subcarriers, modulated with binary phase shift keying (BPSK), quadrature phase shift keying (QPSK), quadrature amplitude modulation (QAM), or 64-QAM to provide different data rates). OFDM allows subcarrier channels to overlap thus providing high spectral efficiency. The modulation technique used by OFDM is more efficient than spread spectrum techniques that are used with IEEE 802.11b; it is the same as is used in 802.11g.

IEEE 802.11a channels

Note For up-to-date information about what channels are supported in your country or region, see the regulatory information related to your country. In addition, not all the clients support all channels.

The following example is for the U.S.-based IEEE 802.11a standard; the 5 GHz unlicensed band covers 300 MHz of spectrum and supports 23 channels. As a result, the 5 GHz band is a conglomeration of three bands in the United States:

- 5.150-to-5.250 GHz (UNII-1)
- 5.250-to-5.350 GHz (UNII-2)
- 5.500-to-5.700 GHz (UNII-2 Extended)
- 5.725-to-5.875 GHz (UNII-3)

IEEE 802.11a operating frequencies and data rates

The IEEE 802.11a standard is resistant to interference from devices that operate in the 2.4 GHz band, such as microwave ovens, cordless phones, and Bluetooth, when it operates in the unlicensed portion of the 5 GHz radio band. Because the IEEE 802.11a standard operates in a different frequency range, it is not compatible with the existing IEEE 802.11b or IEEE 802.11g-compliant wireless devices. But, it does mean that the 2.4-GHz and 5 GHz equipment can operate in the same physical environment without interference.

The IEEE 802.11a standard provides data rates of 6, 9, 12, 18, 24, 36, 48, 54 Mbps, with 54 Mbps being the maximum data rate, though generally at shorter ranges compared to 2.4 GHz network, for a given power and gain. However, it has up to 24 nonoverlapping frequency channels (depending on the geographic area) as compared to the three nonoverlapping channels for the 2.4 GHz band, which results in increased network capacity, improved scalability, and the ability to create microcellular deployments without interference from adjacent cells.

The 5 GHz band in which IEEE 802.11a operates is divided into several sub-bands. Each of the UNII bands presented in the following table were originally intended for different uses, but all can now be used for indoor IEEE 802.11a deployments with applicable power restrictions. Originally, the FCC defined the UNII-1, UNII-2, and UNII-3 bands, each consisting of four channels. The channels are spaced 20 MHz apart with an RF spectrum bandwidth of 20 MHz, thereby providing four nonoverlapping channels.

Table 5: Operating Frequency Range for IEEE 802.11a

Band	Channel ID	Center Frequency
UNII-1	36	5180
	40	5200
	44	5200
	48	5240

Band	Channel ID	Center Frequency
UNII-2	52	5260
	56	5280
	60	5300
	64	5320
	100	5500
	104	5520
	108	5540
	112	5560
	116	5580
	120	5600
	124	5620
	128	5640
	132	5660
	136	5680
	140	5700
UNII-3	149	5745
	153	5765
	157	5785
	161	5805
	165	5825

The limitations on each UNII bands are different such as transmit power, antenna gain, antenna styles, and usage.

- The UNII-1 band is designated for indoor operation, and initially required devices to use permanently attached antennas. The channels in this band (5.150 to 5.250 GHz) are 36, 40, 44, and 48.

- The UNII-2 band is designated for indoor or outdoor operation, and permitted the use of external antennas. The channels in this band (5.250-5.350 GHz) are 52, 56, 60, and 64, and require dynamic frequency selection (DFS) and transmitter power control (TPC).

  Some clients may not support all 5 GHz channels, especially UNII-2 extended channels (100-140). For more information about what channels your country supports, see the regulatory information of your country before you finalize a channel plan. Also note that not all regions support channels 120,124, and 128 (for example United Sates of America and Europe).

- The UNII-3 band, originally intended for outdoor bridge products that use external antennas, is now permitted to be used for indoor or outdoor IEEE 802.11a WLANs as well. The channels in this band are (5.725-5.825 GHz) 149, 153, 157, 161, and 165, and do not require DFS and TPC. Note that not all clients support channel 165.
- The channels in the new frequency range (5.470-5.725 GHz) are 100, 104, 108, 112, 116, 120, 124, 128, 132, 136, and 140, and require DFS and TPC.

Not all channels in a given range can be used in all of the regulatory domains. See the preceding table, which shows the various channels in the UNII-1, -2, and -3 bands, along with the additional 11 new channels.

Note For up-to-date information about what channels are supported in your country or region, see the regulatory information related to your country.

IEEE 802.11a and RToWLAN deployments

While there are as many as 24 nonoverlapping channels in the 5 GHz band, Cisco recommends that you use the lower four and upper four channels of the 5 GHz spectrum as the base for RToWLAN, because they do not have DFS and TPC requirements. Then, determine which other channels are not affected by DFS and TPC, and add these channels to the RToWLAN base of eight channels. The timing requirements of DFS and TPC can adversely affect the RToWLAN call quality. If DFS and TPC can affect the channels in the location that you plan to deploy RToWLAN, you must select the channels appropriately. Otherwise, it is not an issue to select specific channels discretely. Ensure that the channels you select are supported by the WLAN clients (data and RToWLAN). The use of eight non-DFS channels is simpler, but every additional channel that you safely deploy increases the capacity of the design.

In addition to avoiding the DFS and TPC channels, it is also recommended that you avoid adjacent channels in the AP channel layout to avoid interference from the sidebands in each channel. The channel spacing and channel mask characteristics are such that the sidebands produced by an IEEE 802.11a client might interfere with the adjacent channels.

The general power levels and AP separation recommendations that are used in this guide for RToWLAN in the 5 GHz implementation are the same as the 2.4 GHz implementation:

- A power level boundary of ~67 dBm and a separation between adjacent AP channels of -86 dBm.
- A minimum of 20 percent overlap between nonadjacent channels is recommended for 5 GHz bands deployments.
- A 30 percent or higher overlap can still be used for dual-band deployments or for mission-critical environments.

The range in the 5 GHz band is different to that in the 2.4 GHz band. However, when you use the recommended power levels and typical antennas as mentioned in the 5 GHz band example, the obtained distances are similar

to those that are used in the 2.4 GHz example. Therefore, the same AP locations and overlap are used for both 2.4 GHz and 5 GHz bands. The primary difference between the two deployments is the additional capacity that is available due to the additional nonoverlapping channels. This difference is sufficient for the 5 GHz band to be recommended for RToWLAN deployments.

Note The TPC mechanism that is discussed in this section is different from the TPC algorithm that is part of auto-RF.

Single-floor building example

The following figure shows an AP layout that uses eight different channels that is designed to maximize the distance between reused channels. However, in most cases, the requirement is to have more channels available. Because, the 2.4 GHz and 5 GHz AP client radius and co-channel interference radii are the same in this example, the multiple floor examples are not repeated here. The major difference between the two bands is the increase in capacity that is made available by the added channels associated with the 5 GHz band. The more channels that are available for use in the 5 GHz band, the closer the capacity of the system can correlate to the number of APs that are deployed.

Figure 17: 5 GHz Single-Floor Layout

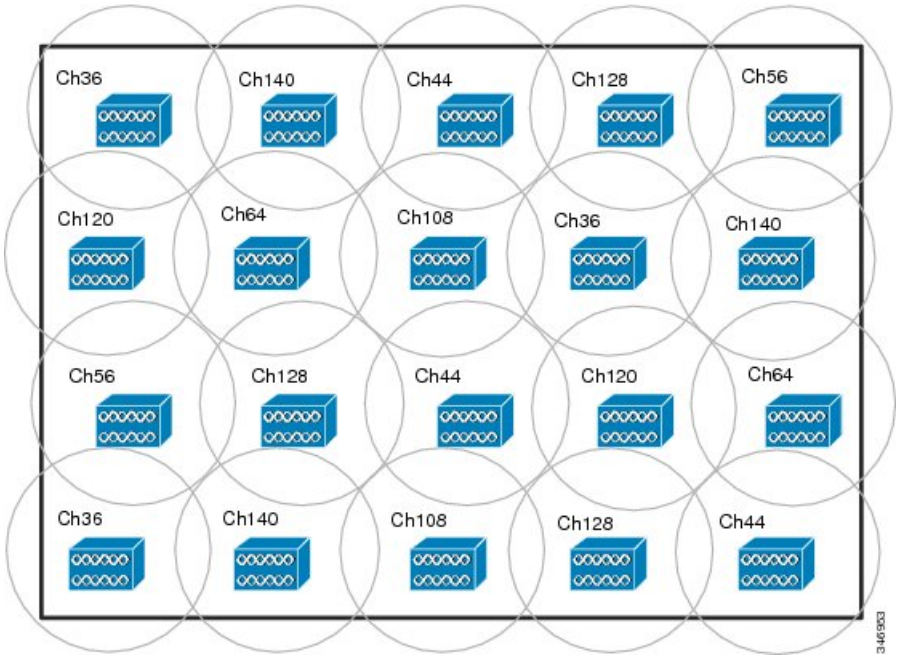

Note The channels used in the preceding figure illustrates the purpose of non-overlapping channels only and is not the recommended channel layouts for any given country or region.

The following figure illustrates an example of the same AP layout as shown in the preceding figure but combined with co-channel interference radius of a single AP.

Figure 18: 5 GHz Single-Floor Layout with Co-channel Radius

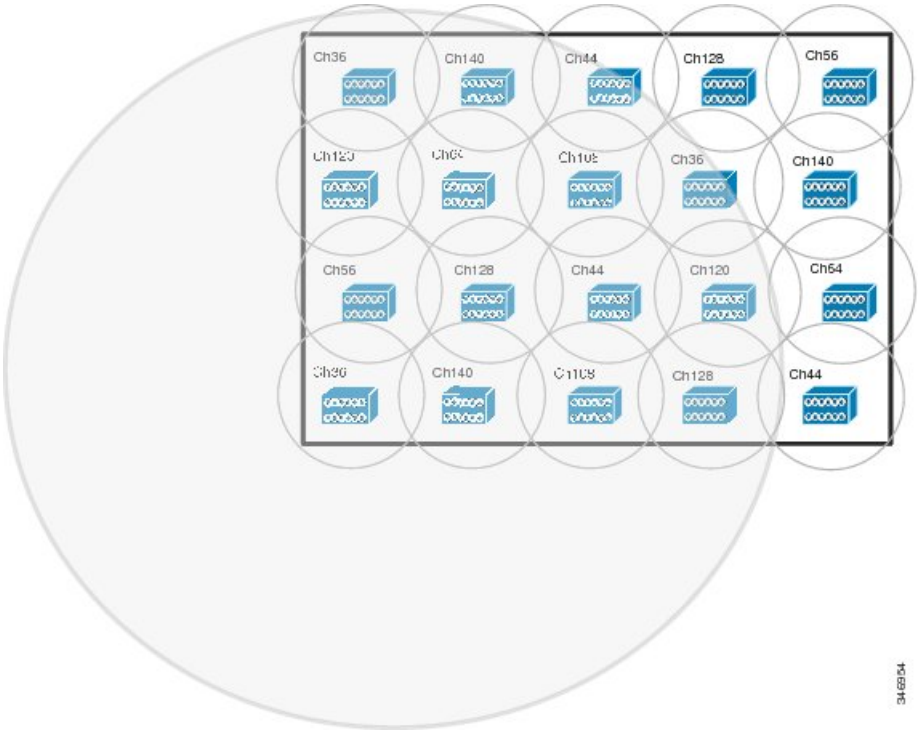

The preceding figure shows that although the co-channel interference is smaller, and more channels are available, the effect of the overall call capacity on the floor is larger. It is difficult to calculate the amount of co-channel interference across the entire floor, given that there are 20 APs and eight channels in use. Therefore, given that there are eight channels in use, the RToWLAN call capacity of the floor is equivalent to eight times the call capacity of a single AP.

Note The channels used in the preceding figure illustrates the purpose of non-overlapping channels only and is not the recommended channel layouts for any given country or region.

Planning tools

The Cisco Unified Wireless Network Wireless Control System (WCS) provides a WLAN planning tool.

The examples that are described in this chapter use simple drawing tools and do not address the complex physical construction and building layout that you must consider for WLAN planning. It is recommended that you use WLAN planning tools to plan the WLAN layout. A small error while planning can prove to be ten times costlier to fix during layout implementation and 100 times costlier to fix during operation. The investments in planning and planning tools to design a WLAN site plan are expensive but contribute to provide maximum benefits after implementation.

Multicast over WLAN networks

Wireless LAN Controller (WLC) configuration options for multicast include changing particular multicast packets to unicast (providing reliable wireless LAN (WLAN) protocol packet delivery), creating defined multicast groups, and prioritizing packets based on application source. WLC configuration plays an important role in maintaining multicast group membership when roaming between access points (APs) or WLCs.

Multicast over WLAN presents delivery issues to Wi-Fi endpoints. These issues are not apparent or common to wired Ethernet where multicast is an effective means of preserving bandwidth on the wired network. Multicast on wireless will in many cases waste the available bandwidth on a WLAN channel if no WLC configuration options are used to manage bandwidth and delivery reliability. Because, multicast streams forwarded to an AP will be forwarded by all the enabled radios on the AP, both the 2.4 GHz and 5 GHz radios on the AP will be forwarding the multicast traffic. With default parameters in place, the multicast traffic is forwarded over the WLAN, even in cases where there are no Wi-Fi endpoints using the multicast application and receiving the multicast traffic. Unnecessary and unused multicast traffic on the Wi-Fi channels impairs the performance of the APs, clients, and the WLAN channels. Additionally, on any client VLAN locally sourced multicast traffic including multicast packets generated by protocols like Hot Standby Router Protocol (HSRP), Protocol Independent Multicast (PIM), Enterprise Interior Gateway Routing Protocol (EIGRP), and Open Shortest Path First (OSPF) will also be flooded throughout the WLAN. All this traffic is sent at the lowest broadcast data rate in use by any client associated to the particular AP potentially reducing throughput on the WLAN. As with multicast traffic on Ethernet, the Wi-Fi endpoint will not acknowledge the receipt of a multicast packet.

A prerequisite for using the multicast performance functionality is that a multicast-enabled network must be configured on all routers between the WLC and the APs. After the administrator enables multicast (multicast mode is disabled by default) and configures a CAPWAP multicast group, the APs download the controller CAPWAP multicast group address during the normal join process (at boot time). Effectively, a CAPWAP multicast group is used to deliver multicast packets to each AP. This allows the routers in the network to use standard multicast techniques to replicate and deliver multicast packets to the APs. For the CAPWAP multicast group, the WLC becomes the multicast source and the APs become the multicast receivers.

Related Topics

20 percent cell overlap calculation
Deployment recommendations for multicast in a wireless network

802.11n and 802.11ac protocols

The design parameters that are covered in this guide and in this chapter are applicable to the newer 802.11n and 802.11ac protocols. A cell edge of -67 dBm is still recommended with these protocols for voice and other jitter sensitive applications. These newer protocols provide greater packet speeds but WLAN usage and application resource demands have increased, so cell capacity planning is still an important aspect of WLAN coverage design.

The 802.11n and 802.11ac protocols, like their predecessors, are half-duplex radio protocols. The major difference is how much frequency is used in the transmission of a data packet. The original 802.11 specification of 1997 defined WLAN channels on 2.4 GHz. The 802.11a specification of 1999 defined the 5 GHz channels. The 802.11n specification remains compatible with both 2.4 GHz and 5 GHz bands. The 802.11n added the concept of high throughput (HT) by adding channel bonding to create cell channels with 40 MHz of frequency. 802.11n also introduced spatial streams and standards-based beamforming. The 802.11ac drops 2.4 GHz from its specification because there is not enough allocated frequency in 2.4 GHz to meet the bandwidth needs of 802.11ac. 802.11n and 802.11ac support the bonding of two or more 20 MHz 802.11 channels to provide a client or an access point with more bandwidth than previous 802.11 protocols.

An 802.11n/802.11ac bonded channel is a WLAN channel that is created by adding the frequency of other 802.11 Wi-Fi channels into a single channel that can become 80 MHz AC Wave1 or 160 MHz AC Wave2. The new 5 GHz protocol 802.11ac is going to have two hardware releases. The first generation of 802.11ac hardware is known as WAVE 1. The next release of 802.11ac hardware is known as WAVE 2. The 802.11ac specification defines a new physical (PHY) specification that provides very high throughput (VHT) based on the orthogonal frequency division multiplexing (OFDM) radio modulation system. The specification increases the number of spatial streams and provides for multi-user transmissions. WAVE 2 will provide the bonding of two 80 MHz channel into a 160 MHz channel. The 160 MHz channel can be contiguous channels or two 80+80 MHz non-contiguous channels.

802.11ac introduced support of multi-user multiple input – multiple output (MIMO) (multiple clients with spatial stream support) to transmit and receive unique data streams simultaneously. These specification enhancements have increased the throughput and bandwidth in the coverage area of the access point. Beamforming increases the AP to an individual clients performance and also increases the bandwidth on that AP WLAN channel. The MIMO antenna support at the client radio or AP radio improves transmission and reception quality. These technology improvements increase capacities in a coverage area. As with the legacy technologies of the 1990s, the AP coverage is still an important design consideration for the application performance of Wi-Fi endpoints.

802.11ac VHT PHY provides backward compatibility to 802.11n and 802.11a protocols. Therefore, a legacy 802.11a client can associate, authenticate, and pass traffic through the 802.11ac AP at the data rates that are supported on the client. Current 802.11ac APs support 4x4 MIMO technology with three spatial streams that provides transmission data rates of 1.3 Gbps.

Client device application performance in the new VHT client radio hardware is directly related to the current bandwidth of the WLAN channel. This was a condition with the legacy clients and will be a condition with the next generation of Wi-Fi protocols. Therefore, the same design criteria that is used for earlier WLAN cell design still applies. WLAN applications generally require more bandwidth, as much bandwidth as wired applications. To achieve more bandwidth in a given floor space, WLAN channels in the given floor spaces must become more efficient. This requires more efficiency on the client radio and the AP radio. This may also require configuration changes of the data rates that are assigned to the radios. Removing the 1997 data rates of 1 Mbps and 2 Mbps is highly recommended. Removing the 1999 data rates of 5.5 Mbps and 11 Mbps is also recommended when applicable. Having any of those data rates set as required is likely to impact client application performance in dense coverage area. Enabling 5 GHz is highly recommended.

Successfully deploying a WLAN network that is able to provide diverse services in a challenging environment can be a challenging project for any organization. Doing it right the first time requires a special set of skills and knowledge that sometimes is difficult to find. Partnering with a network integrator that is experienced with wireless deployments can be very beneficial.

The 802.11ac specification is supported by the AP3600 with and 802.11ac module and the AP3700. Current radio technology supports 80 MHz wide channels but not multi-user transmissions or 160 MHz wide channels.

Related Topics

Data Sheets for Cisco APs supporting 802.11ac
802.11ac white paper

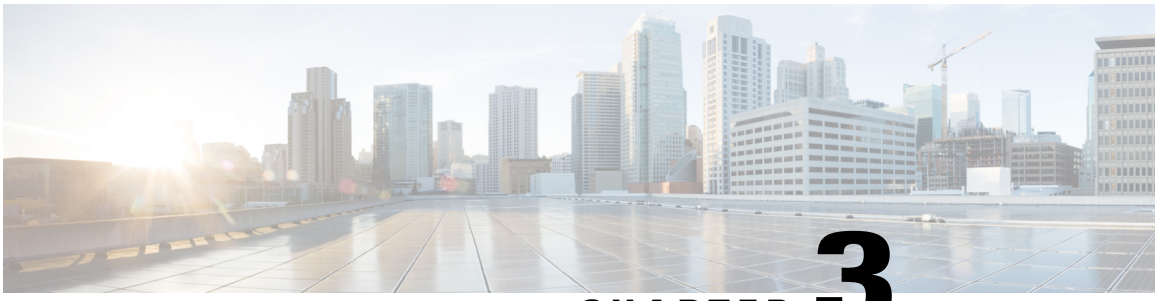

CHAPTER 3

Real-Time Traffic over WLAN Quality of Service

This chapter describes quality of service (QoS) in Real-Time Traffic WLAN implementations. It describes WLAN QoS in general and does not discuss security and segmentation in depth, though QoS is a part of these components. It also provides information about Cisco Centralized WLAN architecture features.

Quality of Service architectural overview

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

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

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

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

QoS importance to Real-Time Traffic over WLAN

The WLAN used for both packet transmission and reception is unlicensed, unprotected, and unshielded. Multiple specifications, protocols, and devices take advantage of unlicensed and no-cost media (radio frequencies) by a WLAN. Consider the following example.

A tablet user in a business office is using Bluetooth to print a document. Another laptop user in the same office is using 2.4 GHz frequency Wi-Fi for a video conference and presentation. A new guest user, in the lobby, is using a smartphone to check email through the guest VLAN on the Wi-Fi network. The Wi-Fi network must prioritize the 2.4 GHz radio frequency shared by the three devices, to give real-time video conference application priority over the guest smartphone user and tablet user. In addition, the Wi-Fi network must also address the tablet Bluetooth transmission interference.

WLAN queue and schedule mechanism

802.11 WLAN has its own queue and schedule mechanism, which is divided into four access categories (ACs). These four Wi-Fi AC queues provide differentiated access to the Wi-Fi channel. The four Wi-Fi QoS categories also align to the 802.1P access categories. Because Wi-Fi is designed to carry multiple Layer 3 protocols, specific implementations typically map the Differentiated Services Code Point (DSCP) values in the IP header of the packet to be sent by the Wi-Fi radio into one of the ACs.

The voice packets are placed in queue with the highest priority for WLAN depending on the DSCP value or IP QoS value of voice packets in the phone call, known as voice access category. In Cisco wireless controllers, voice packets also map to the Platinum QoS profile. Voice and video packets from a voice or video call have quicker and more frequent access to the Wi-Fi channel than data packets. There will be packet collisions between the phone call and data application, because Wi-Fi is a shared medium. Wi-Fi QoS prioritizes the backoff and packet retry logic for both real-time voice and video traffic and data traffic, based on configuration values in the WLAN enhanced distributed coordination function (EDCF).

The Bluetooth (BT) radio, as well as the Wi-Fi radios in the laptop, smartphone, tablet, and access point are all half-duplex. Therefore, when each of these four radios transmit a packet, they all change to a receive-packet state, waiting for packet that acknowledges that the packet they sent was received correctly. This is where the 802.11e specification of 2005 for Wi-Fi plays an important role with QoS channel prioritization. 802.11 WLAN protocols have at their basic media access logic a process called carrier sense multiple access with collision avoidance (CSMA/CA). The Wi-Fi units wait in receive packet mode for the absence of a carrier (radio frequency) before they transmit a packet. Therefore, all Wi-Fi devices in the vicinity of a BT radio wait for the BT radio to quit transmitting before they can transmit. The dominant BT data rate used in smartphones to earbuds is 2 Mbps. Therefore, a G.711 voice packet of 256 bytes sent from a BT radio is going to delay the Wi-Fi devices on the same 2.4 GHz frequency for over 1100 microseconds, while the same G.711 voice packet on 802.11n Wi-Fi takes about 50 microseconds to send.

The Cisco CleanAir technology defines and locates the general area where BT and other interferers exist and helps avoid them. But, BT uses the entire 2.4 GHz frequency allocation that Wi-Fi uses. Therefore, it is a

protocol that cannot be avoided by all the Wi-Fi channels thus making the QoS mechanism the best solution for mitigating BT interference.

The Wi-Fi QoS protocol is known as Wi-Fi Multimedia (WMM). WMM is a subset of the 802.11e specification. The 802.11e specification was approved in 2005; however, it was extensively used by Wi-Fi Alliance and Microsoft even before 2005. With the 802.11e specification, devices required new drivers to become Wi-Fi QoS capable with no hardware changes. The legacy devices without QoS have specialized hardware designs with limited firmware memory.

With the approved Wi-Fi QoS protocol in 2005, the need for Wi-Fi channel bandwidth has grown immensely along with continuous improvement in quality, range, and speed. Thus, it is now possible for a site that was dependent on the legacy data rates of 1, 2, 5.5 and 11 Mbps to disable those rates completely with valid reasons.

802.11 WLAN channel bandwidth is a managed media resource. The Wi-Fi channel is the first hop upstream and the last hop downstream. Because this medium is open to radio interference and non-Wi-Fi protocols, it is the media hop that influences the performance of the applications running on the Wi-Fi devices. When considering voice and video applications like Cisco Jabber, the Wi-Fi channel is the medium most likely to have a negative impact on the mean opinion score (MOS) value of a call. Therefore, bandwidth must be managed to insure that applications can perform to meet the expectations of the users.

Removing 802.11b data rates can double bandwidth in a 2.4 GHz Wi-Fi channel. The sites that require 802.11b and have performance issues must consider whether to disable the data rates of 1, 2 and 5.5. A required data rate of 11 Mbps provides all the support required for legacy and special application devices. Cisco APs with beam forming technology use the orthogonal frequency-division multiplexing (OFDM) modulation to further enhance client performance and Wi-Fi channel bandwidth. The OFDM radio modulation was introduced to Wi-Fi in 1999 in the 5 GHz bands with the 802.11a specification. OFDM came into practice for 2.4 GHz with the 802.11g specification that was approved in 2003. With over 10 years of use and technology development, 2.4 GHz OFDM modulation eliminates the high cost of maintaining backwards compatibility for 802.11b technology.

The following four WMM QoS priority options, also known as QoS profiles, are used for WLAN/SSID configuration with Cisco Wireless LAN Controllers:

- Platinum
- Gold
- Silver
- Bronze

These WMM options set the priority limit of traffic between the AP, the wireless LAN controller upstream, and the priority limit from the wireless controller to the Wi-Fi endpoint downstream.

For example, a voice packet with a DSCP voice priority value in a WLAN/SSID configuration of Silver with best effort has an IP Control and Provisioning of Wireless Access Points (CAPWAP) wrapper header DSCP value of best effort. A data packet with a DSCP value of zero in a WLAN/SSID with a Platinum voice configuration maintains a best effort priority, while a voice packet in the same WLAN/SSID maintains its voice priority.

The Cisco WLAN Controller (WLC) provides several methods to upgrade video or voice packets that originate with a DSCP value less than a video or voice value to be upgraded to the highest priority that is configured for the WLAN/SSID, for example, a desktop client like Cisco Jabber on a Windows computer. While the call control server can be configured to have the software application mark audio and video packets with appropriate DSCP values, the Windows operating system may be in a default configuration that does not allow QoS marking. Therefore, the audio and video is sent to the AP with best effort Wi-Fi media access. But upstream

from the AP, those packets can be inspected using deep packet snooping, resulting in subsequent packet marking upgrades.

The Cisco Jabber application with Cisco Enterprise Medianet has the advantage of the peer relationship with Medianet switches. A Medianet-aware client application has built-in intelligence to address the unique challenges of video and rich media by providing:

- Enhanced video performance and end-user quality of experience (QoE) over the network
- Simplified installation and management of video endpoints
- Faster troubleshooting for voice, data, and video applications
- Capability to assess effect of video, voice, and data in the network

The first hop to the network from a Wi-Fi client is the shared Wi-Fi channel. The client device access to the Wi-Fi channel is the most influential aspect in the overall user acceptance of performance.

The WLAN configuration holds no control on the marking of the packets from the endpoint Wi-Fi client to the AP or vice-versa. The application DSCP markings, operating system, and the WMM driver control the marking values and the AC queues. Therefore, is it most important that you manage these three aspects of the source client. The delay incurred on the first hop because of the lack of QoS on the endpoint Wi-Fi client cannot be made up by deep packet inspection and remarking logic in the AP or upstream from the AP.

Wireless QoS deployment schemes

Cisco Unified Wireless products support WMM, a QoS system based on IEEE 802.11e published by the Wi-Fi Alliance, WMM Power Save, and Admission Control.

The following figure shows an example of deployment of wireless QoS based on the Cisco Unified Wireless technology features.

Figure 19: QoS Deployment Example

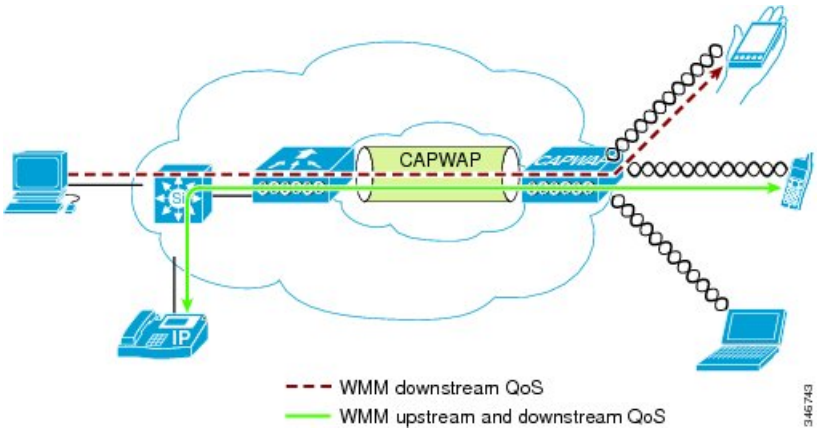

QoS parameters

QoS is the measure of performance for a transmission system that reflects its transmission quality and service availability. Service availability is a crucial element of QoS. Before you implement QoS successfully, the

network infrastructure must be highly available. The network transmission quality is determined by latency, jitter, and loss, as described in the following table.

Table 6: QoS Parameters

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

Upstream and downstream QoS

The following figure defines radio upstream and downstream and includes the following:

- **Radio downstream QoS:** Traffic leaving the AP and traveling to the WLAN clients. Radio downstream QoS is the most common deployment. The radio client upstream QoS depends on the client implementation.
- **Radio upstream QoS:** Traffic leaving the WLAN clients and traveling to the AP. WMM provides upstream QoS for WLAN clients that support WMM.
- **Network downstream:** Traffic leaving the WLC and traveling to the AP. QoS can be applied at this point to prioritize and rate-limit traffic to the AP. Configuration of Ethernet downstream QoS is not covered in this chapter.

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

Figure 20: Upstream and Downstream QoS

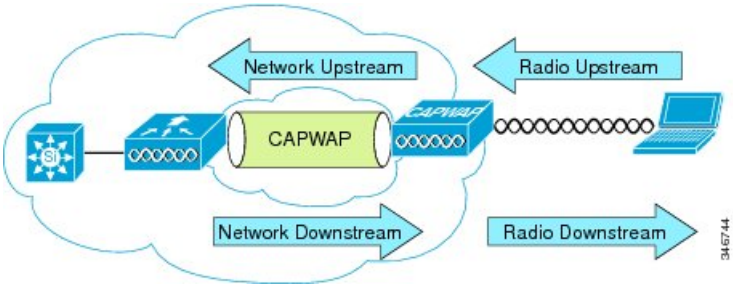

QoS/WMM and Wi-Fi channel/network performance

The application of QoS features cannot be detected on a lightly loaded network. QoS features begin to apply on the application performance as the load on the network increases. If you can measure latency, jitter, and loss when the medium is lightly loaded, it indicates either a system fault, poor network design, or that the latency, jitter, and loss requirements of the application are not a good match for the network.

QoS functions to keep latency, jitter, and loss for selected traffic types within acceptable boundaries. When you provide only radio downstream QoS from the AP, radio upstream client traffic is treated as best-effort. A client must compete with other clients for upstream transmission, as well as compete with best-effort transmission from the AP. At certain load conditions, a client experiences upstream congestion, and the performance of QoS-sensitive applications becomes unacceptable, despite the QoS features on the AP. The upstream and downstream QoS can be operated either by using WMM on both the AP and WLAN client, or by using WMM and a client proprietary implementation.

Note Benefits to client traffic from a WLAN client with WMM support are not direct. The applications that look for benefits of WMM assign a priority classification to their traffic, and the operating system passes this classification to the WLAN interface. In purpose-built devices, such as wireless voice handsets, the implementation is integrated as part of the design. However, if you implement it on a general-purpose platform, such as a personal computer (PC), you must first implement application traffic classification and OS support to achieve better results.

Wi-Fi multimedia

Wi-Fi MultiMedia (WMM), formerly known as Wireless Multimedia Extensions, refers to QoS over Wi-Fi. QoS enables Wi-Fi access points to prioritize traffic and optimize the way shared network resources are allocated among different applications.

This section describes the following three considerations for WMM implementation:

- WMM access
- WMM classification
- WMM queues

WMM access

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

WMM classification

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

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

Table 7: 802.1P and WMM Classification

Priority	802.1P Priority	802.1P Designation	Access Category	WMM Designation
Lowest	1	BK	AC_BK	Background
	2	-		
	0	BE	AC_BE	Best Effort
	3	EE		
	4	CL	AC_VI	Video
	5	VI		
	6	VO	AC_VO	Voice
Highest	7	NC		

The following figure shows the WMM data frame format.

Figure 21: WMM Frame Format

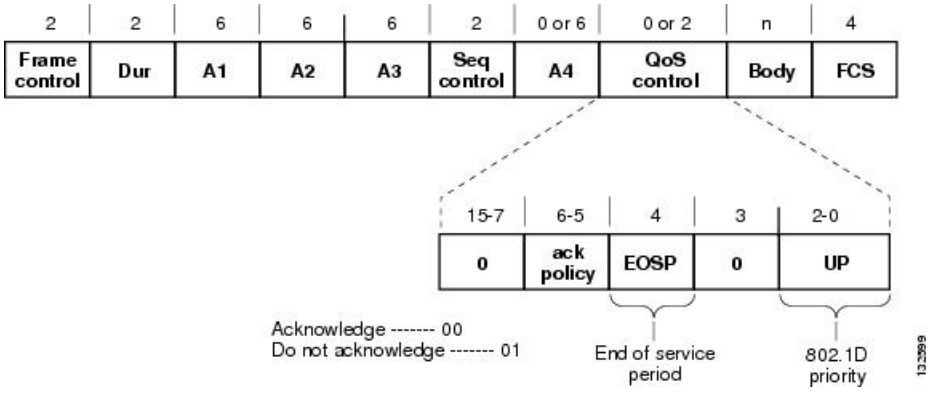

Even though WMM maps the eight 802.1P classifications to four access categories, the 802.1D classification is sent in the frame.

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

WMM queues

The following figure shows the queuing that is performed on a WMM client or AP.

Figure 22: WMM Queues

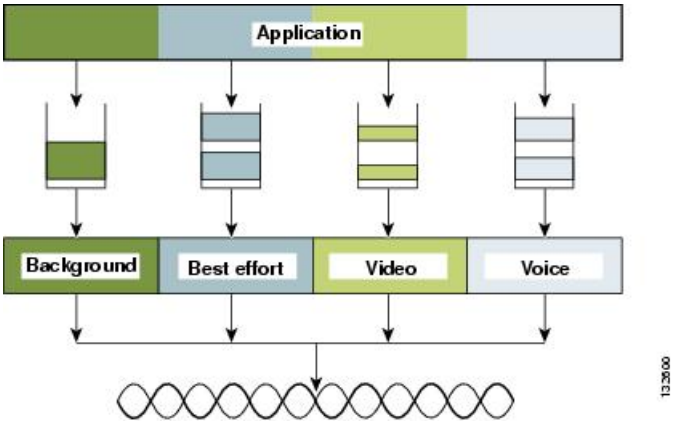

There are four separate queues, one for each of the access categories. Each of these queues compete for the wireless channel, with each of the queues using different interframe space, contention window (CW) minumum (CWmin) and contention window maximum (CWmax) values as defined by EDCF. If more than one frame from different access categories collide internally, the frame with the higher priority is sent, and the lower priority frame adjusts its backoff parameters as though it had collided with a frame external to the queuing mechanism.

The following figure shows the principle behind EDCF where different interframe spacing and CWmin and CwMax values (for clarity, CwMax is not shown) are applied per traffic classification.

Figure 23: Access Category (AC) Timing

Different traffic types can wait different interface spaces before counting down their random backoff, and the CW value used to generate the random backoff number also depends on the traffic classification. For example, the CWmin[3] for voice traffic is 2^3-1, and CWmin[5] for best effort traffic is 2^5-1. High priority traffic has a small interframe space and a small CWmin value, giving a short random backoff, whereas best effort traffic has a longer interframe space and large CWmin value, that, on average, gives a large random backoff number.

The following figure shows the WMM information in a probe response.

Figure 24: Probe Response WMM Element Information

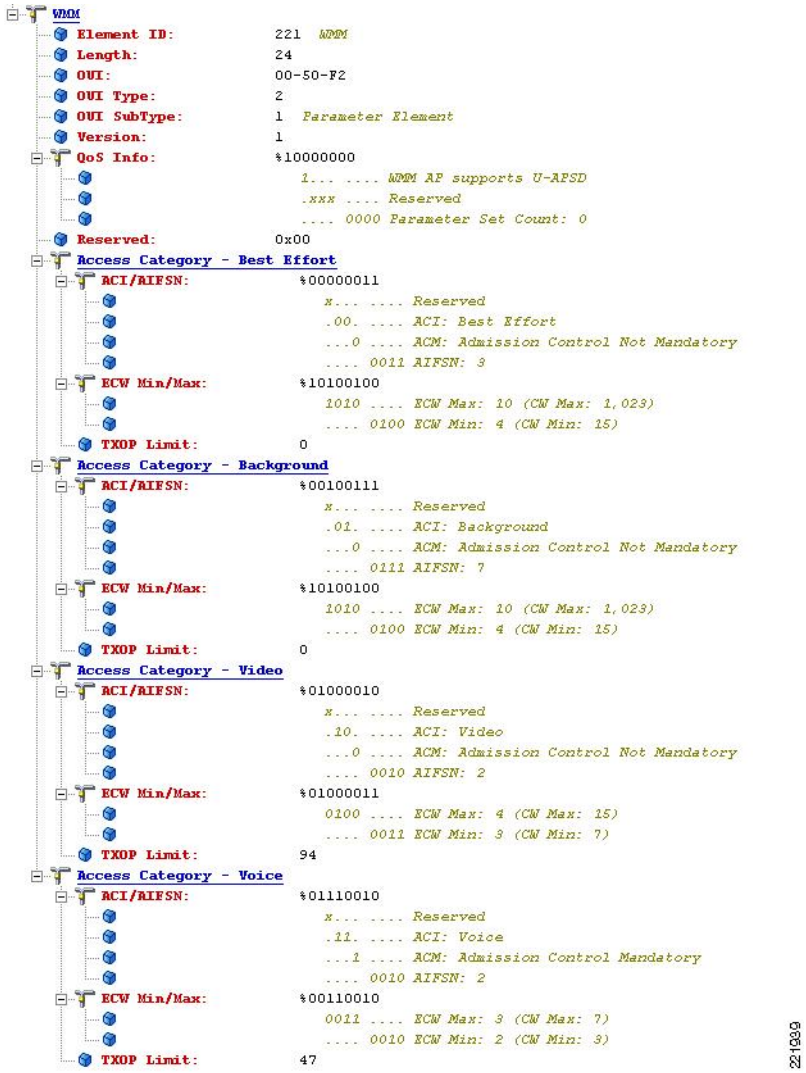

The elements on the client not only contain WMM AC information, but also define which WMM categories require admission control. For example, in the preceding figure, the admission control for voice AC is set to **Mandatory**. Therefore, the client is required to send the request to the AP, and have the request accepted, before it can use this AC.

Unscheduled automatic power-save delivery

Unscheduled automatic power-save delivery (U-APSD), a WMM feature of Wi-fi devices, provides two key benefits:

- Allows the voice client to synchronize the transmission and reception of voice frames with the AP, allowing the client to transition into power-save mode between the transmission or reception of each voice frame tuple.

  The WLAN client frame transmission in the access categories supporting U-APSD triggers the AP to send any data frames that are queued for that WLAN client in that AC. A U-APSD client remains listening to the AP until it receives a frame from the AP with an end-of-service period (EOSP) bit set. Once the client receives a frame with the EOSP bit set which indicates there are no other frames, the client goes back into power-save mode. This triggering mechanism is a more efficient use of client power than the regular listening for beacons method, at a period controlled by the delivery traffic indication map (DTIM) interval. This is because the latency and jitter requirements of voice and video are such that a voice and video over IP (VVoIP) client would either not be in power-save mode during a call, resulting in reduced talk times, or would use a short DTIM interval, resulting in reduced standby times.

  U-APSD allows the use of long DTIM intervals to maximize standby time without sacrificing call quality. You can apply this feature individually across access categories; however, only voice ACs in the AP use U-APSD and other ACs still use the standard power-save feature.

- Increases call capacity

  The coupling of transmission buffered data frames from the AP with the triggering data frame from the WLAN client allows the frames from the AP to be sent without the accompanying interframe spacing and random backoff, thereby reducing the network contention.

The following figure shows an example of traffic flow with U-APSD.

Figure 25: U-APSD Traffic Flow

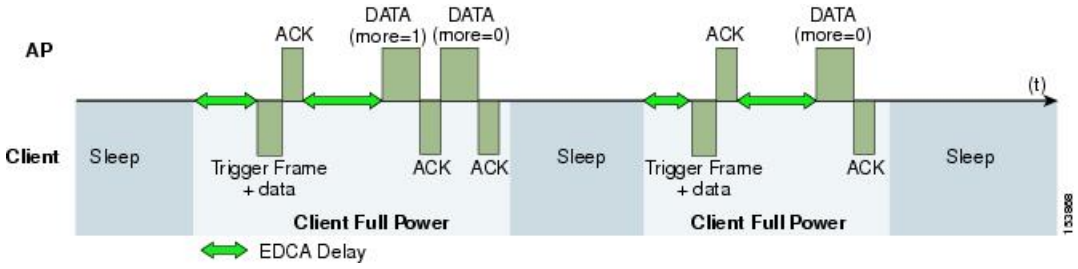

In this example, the trigger for retrieving traffic is the client sending traffic to the AP. When the AP acknowledges the frame, it indicates the client that data is in queue and must wait. The AP then sends data to the client typically as a transmit opportunity (TXOP) burst where only the first frame has the EDCF access delay. All subsequent frames are then sent directly after the acknowledgment frame. In a Real-Time Traffic over WLAN implementation, only one frame is queued at the AP, and the real-time-capable WLAN client becomes idle after receiving that frame from the AP.

The U-APSD approach overcomes both the disadvantages of the previous scheme, thus making it efficient. The timing of the polling is controlled through the client traffic, which in the case of voice and video is symmetric. If the client is sending a frame every 20 ms, it waits to receive a frame at each 20 ms time interval. This introduces a maximum jitter of 20 ms, rather than n * 100 ms jitter.

The following figure shows an example frame exchange for the standard 802.11 power-save delivery process.

Figure 26: Standard Client Power-Save

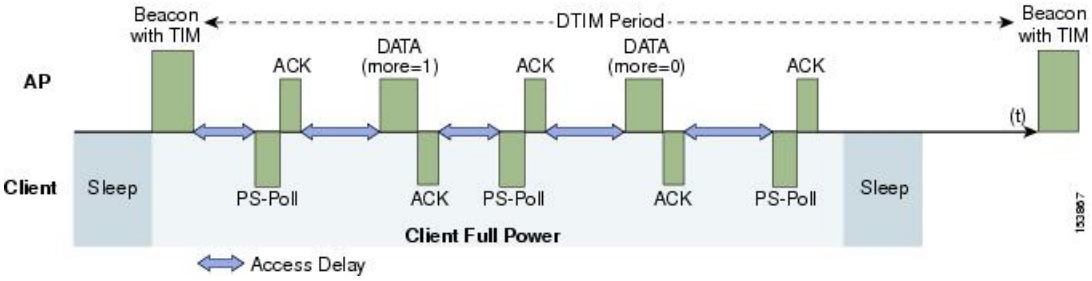

The client in power-save mode first detects that there is data waiting for it at the AP from the traffic indicator map (TIM) in the AP beacon. The client must power-save poll (PS-Poll) the AP to retrieve that data. If the data that is sent to the client requires more than one frame to be sent, the AP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

The standard client power-save has two disadvantages.

- It is inefficient for the PS polls and the normal data exchange to go through the standard access delays associated with distributed coordination function (DCF).
- Retrieving the buffered data is dependent on the DTIM, which is an integer multiple of the beacon interval. Standard beacon intervals are 100 ms. This introduces a level of jitter that is unacceptable for voice and video calls, and voice and video capable wireless endpoints handsets switch from power-save mode to full transmit and receive operation when a call is in progress.

This standard client power-save mode gives acceptable voice and video quality but reduces battery life. The Cisco Unified Wireless IP Phones address this issue by providing a PS-Poll feature that allows the phone to generate PS-Poll requests without waiting for a beacon TIM. This allows the device to poll for frames when it has sent a frame, and then go back to power-save mode. This feature does not provide the same efficiency as U-APSD, but improves battery life for Cisco Unified Wireless IP Phones on WLANs without U-APSD.

Traffic Specification Admission Control

Traffic Specification (TSPEC) allows an 802.11e client to signal its traffic requirements to the AP. In the 802.11e media access control (MAC) definition, the following two mechanisms provide prioritized access, both provided by the transmit opportunity (TXOP):

- Contention-based EDCF option
- Controlled access option

With the TSPEC features, a client can specify its traffic characteristics, which automatically results in the use of controlled access mechanism. The controlled access mechanism enables the client to grant a specific TXOP to match the TSPEC request. However, the reverse mechanism is also possible; that is, a TSPEC request can be used to control the use of various ACs in EDCF. In a TSPEC mechanism, a client must send the TSPEC request before it sends any priority-type traffic.

For example, a WLAN client device that requires to use the voice AC must first make a request for use of that AC. You can configure the use of voice and video ACs by TSPEC requests but the use of best effort and background ACs can happen without TSPEC requests.

The use of EDCF ACs, rather than the 802.11e hybrid coordinated channel access (HCCA), to meet TSPEC requests is possible because the traffic parameters are simple to allow them to be met by allocating capacity, rather than creating a specific TXOP to meet the application requirements.

Add traffic stream

The Add traffic stream (ADDTS) function is how a WLAN client performs an admission request to an AP. Signaling its TSPEC request to the AP, the admission request can be in two forms:

- **ADDTS action frame:** Used when a voice or video call is originated or terminated by a client associated to the AP. The ADDTS contains TSPEC and might contain a traffic stream rate set (TSRS) information element (IE) (Cisco Compatible Extensions Version 4 clients).
- **Re-association message:** Uses the re-association message when the re-association message contains one or more TSPEC and one TSRS IE if an STA roams to another AP.

The TSPEC element in ADDTS describes the traffic request. Apart from data rates and frame sizes, the TSPEC element also tells the AP the minimum physical rate that the client device will use. This helps to determine the time that the station consumes to send and receive in this TSPEC, therefore allowing the AP to calculate whether it has the resources to meet the TSPEC. The WLAN client (VoIP handsets) uses TSPEC admission control during a call initiation and roaming request. While the WLAN client is roaming, the TSPEC request is appended to the reassociation request.

Related Topics

Enterprise Mobility Design Guide

Client connection types

The following figure shows `Monitor` > `Clients` page of the Cisco WLAN Controller (WLC) which indicates what Wi-Fi protocol the client used to associate to the WLAN. In the figure below, the client is connected to the 2.4 GHz channel because the protocol is 802.11bn. The client can also be 802.11b, 802.11g or 802.11n. The client is connected to the 5 GHz channel if the protocol is 802.11an. Clicking on the MAC Address link shows the characteristics of the selected client.

Figure 27: WLAN Controller Clients

The following figure shows `Clients` > `Detail` page, which displays detailed client information available on the WLC. This page displays three important fields and the values about the client connection status:

- **Current Tx-Rate-Set:** Indicates the data rate, here m15.

- **RSSI:** A value of -39 dBm indicates a strong signal.

  For real-time traffic applications, the desired receive signal strength indicator (RSSI) is a strength of -67 dBm at the cell edge.

- **QoS Level:** Set to Platinum, indicates the client can send and receive at the highest WMM priority.

Figure 28: WLAN Controller Clients Detail Page 1 of 2

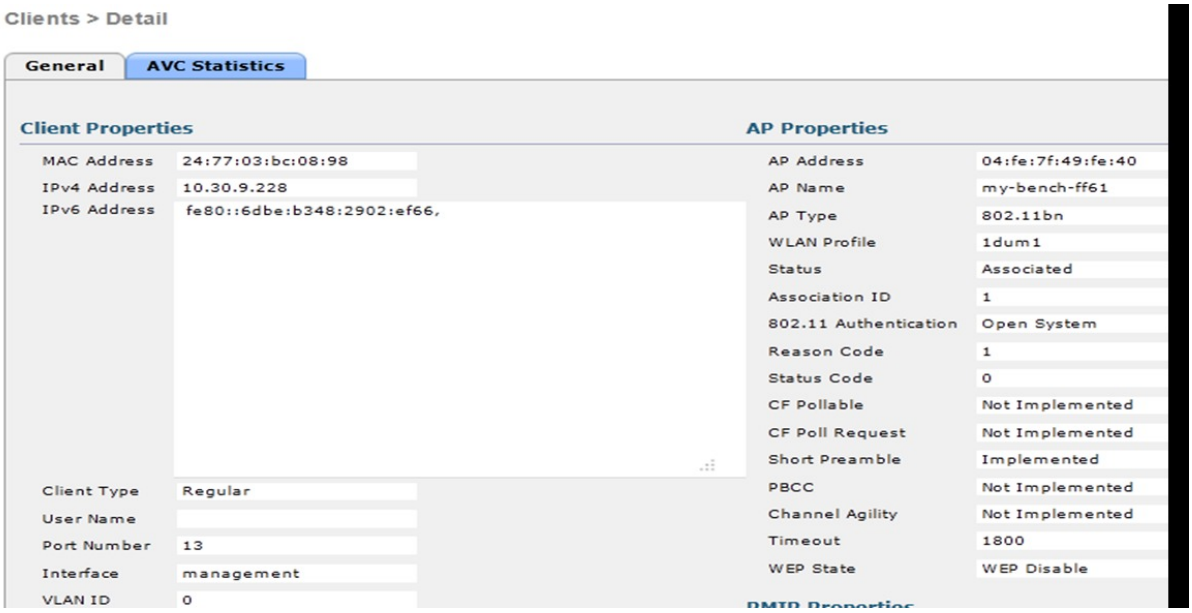

Figure 29: WLAN Controller Clients Detail Page 2 of 2

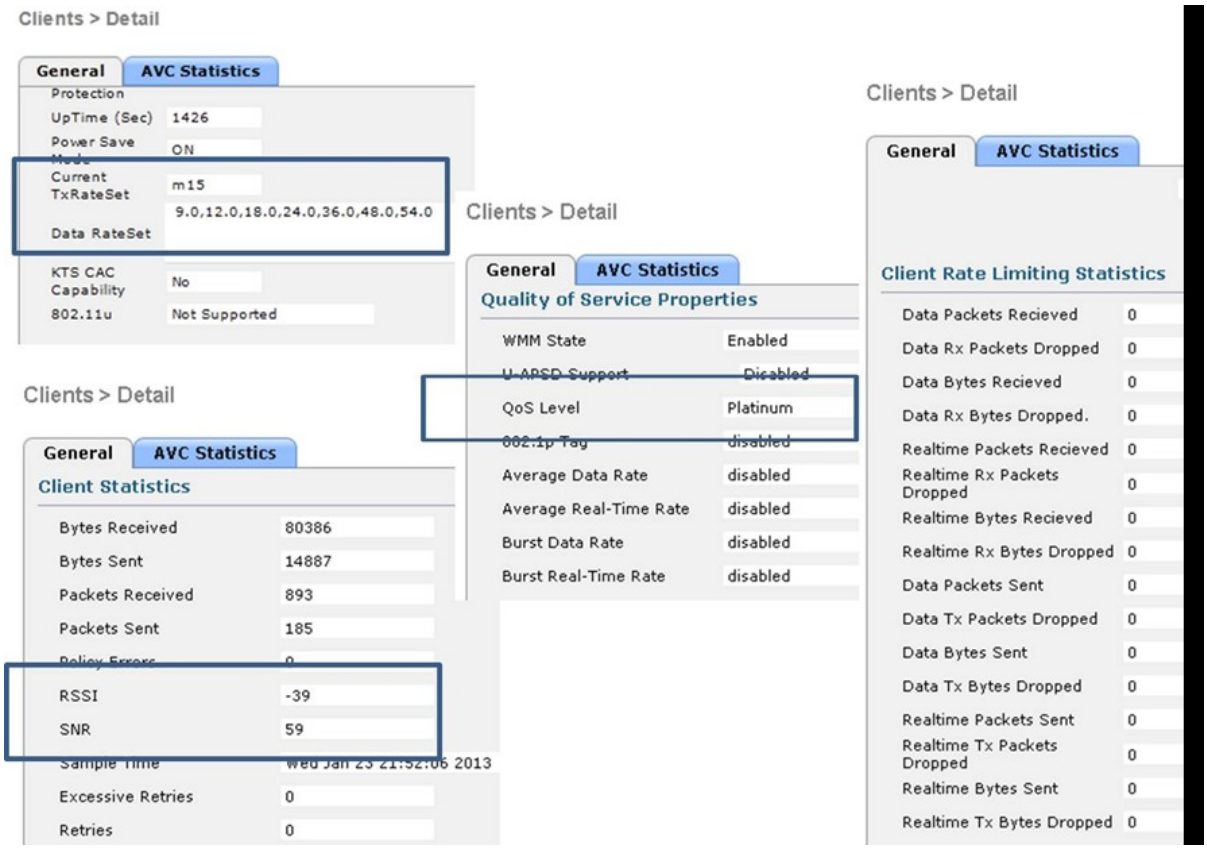

The values in the rate limiting column indicate this client is not part of a rate limiting profile.

The following figure shows the load statistics of the AP that is associated to the client.

Figure 30: WLAN Controller Channel Utilization

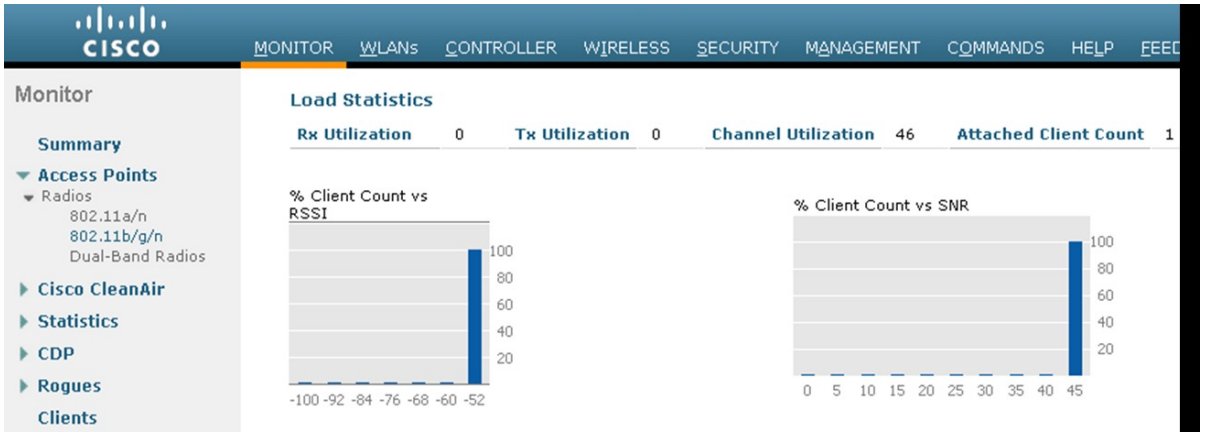

The AP radio is 802.11b/g/n and the channel utilization for the client and the AP is 46 percent. The client and the AP utilization are both 0 percent because they are not sending a significant number of packets to each other. However, the channel that the client and AP are using is very busy with traffic from other APs, other clients, and interference.

The 46 percent channel utilization is above the channel utilization wireless packetized ALOHA standard. The ALOHA protocol defines a radio channel as full when channel utilization reaches 33 percent. This means that the channel is busy, so the packets must wait for an open time slot before they are transmitted. This level of channel utilization is not uncommon with Wi-Fi 2.4 GHz channels. In this scenario, QoS helps to manage channel bandwidth. This is also the reason for Wi-Fi call admission control (CAC). CAC is a part of the 802.11e specification.

The following figure shows WLC CAC configuration page.

Figure 31: WLC Call Admission Control Settings

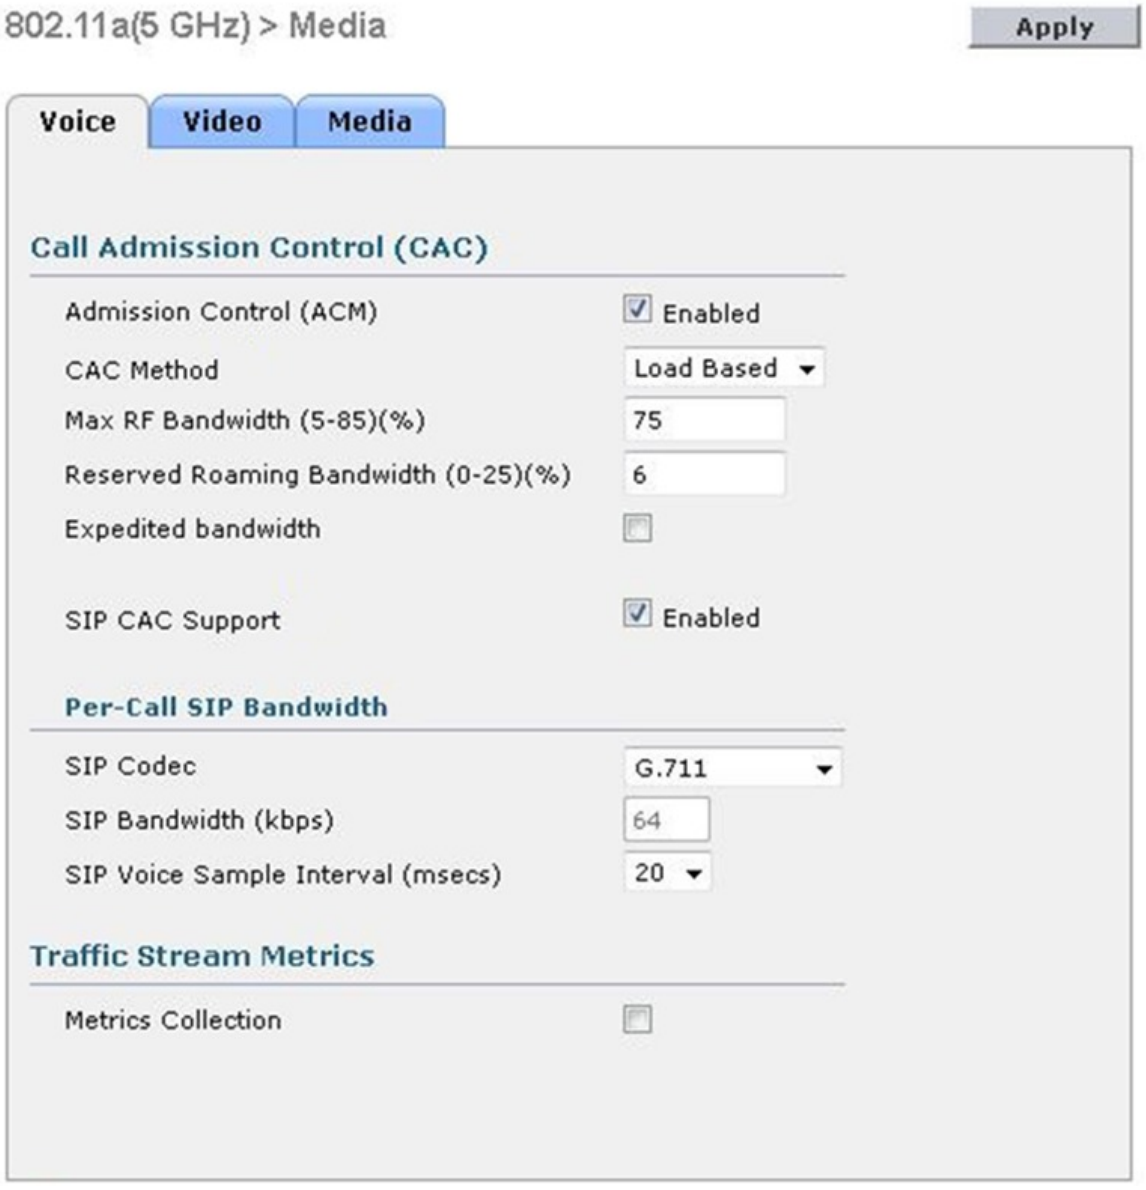

Admission control mandatory (ACM) Load Based CAC for wireless phones and other devices is effective to maintain good quality calls and preserve bandwidth. Load-based CAC measures the load of the Wi-Fi which

is best in high density deployments where there is a high level of channel reuse across several APs. Cisco also supports SIP CAC. For SIP CAC, the WLAN must have media session snooping enabled. If the CAC method is load-based, then SIP CAC also uses channel load. Most softphones and smartphones use SIP as the call connection protocol, so SIP CAC is important. When you enable SIP CAC and deploy TCP-based SIP clients, in scenarios where there is not enough bandwidth for a new voice or video call to go through, the WLAN network stops forwarding SIP frames upstream and downstream. Based on client code behavior, this may cause loss of call control registration. In the case of SIP CAC with UDP-based SIP clients, the WLAN network will send a 486 Network Busy message. Based on the client code behavior, client may roam to another AP or terminate call setup. In addition to CAC configuration for voice traffic, there are tabs for video and media traffic. These provide configuration options to extend CAC to video and media. With the help of these tabs, you can configure how the bandwidth of a Wi-Fi channel is divided between real-time voice and video applications and media applications, which in turn determines how much bandwidth remains for data applications.

For SIP-based Cisco WLAN endpoint and mobile client deployments, Cisco recommends not to enable SIP CAC support because they utilize TCP-based SIP versus UDP-based SIP.

Related Topics

ALOHA mobile network protocol definition

Cisco Wireless LAN Controller Configuration for CAC configuration based on model of WLC

QoS advanced features for WLAN infrastructure

The Cisco Centralized WLAN architecture has multiple QoS features, in addition to WMM support. Cisco WLAN Controller QoS profiles are the primary mechanism for implementing advanced QoS feature. The following four QoS profiles and corresponding traffic types are supported:

- Platinum: voice application traffic
- Gold: video application traffic
- Silver: best effort traffic
- Bronze: background traffic

The following figure shows the four available QoS profiles on the Cisco WLAN Controller.

Figure 32: WLAN Controller QoS Profiles

For each profile, you can configure the bandwidth contracts, RF usage control, and the maximum IEEE 802.1P classification that is allowed.

Figure 33: Editing WLAN Controller QoS Profiles

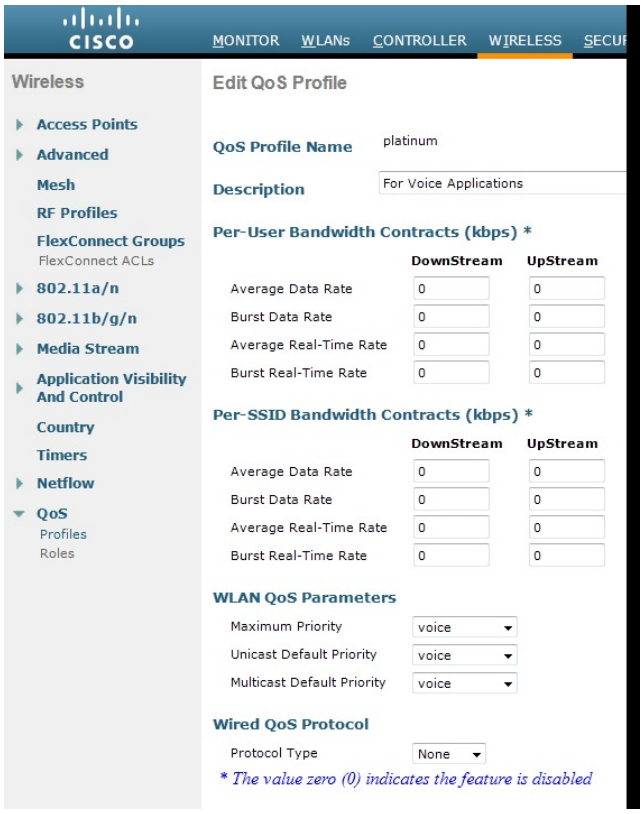

Cisco recommends that you use the default values for **Per-User Bandwidth Contracts** settings and use IEEE 802.11 WMM features to provide differentiated services.

For WLANs that use a given profile, the IEEE 802.1P classification in that profile controls two important behaviors:

- Determines what class of service (CoS) value is used for packets that are initiated from the WLC.

  The CoS value that is set in the profile is used to mark the CoS of all CAPWAP packets for WLAN using that profile. So, for a WLAN with platinum QoS profile, and the IEEE 802.1P mark of 6, will have its CAPWAP packets from the AP Manager interface of the controller marked with CoS of 5. The controller adjusts the CoS to be compliant with Cisco QoS baseline recommendations. If the network is set to trust CoS rather than a DSCP at the network connection to the WLC, the CoS value determines the DSCP of the CAPWAP packets that are received by the AP, and eventually the WMM classification and queuing for WLAN traffic, because the WLAN WMM classification of a frame is derived from the DSCP value of the CAPWAP packet carrying that frame.

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

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

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

The key point in Unified Wireless Network is that you must always consider IEEE 802.11e classifications, and allow the Unified Wireless Network Solution to take responsibility to convert between IEEE classification and the Cisco QoS baseline.

For more information about Per-User Bandwidth Contracts, Per-SSID Bandwidth Contracts, and WLAN QoS Parameters, see the WLC configuration guides that match the WLC code release and model.

You can configure WLAN with various default QoS profiles as shown in the following figure.

Figure 34: WLAN Controller WLAN Default QoS Profile Settings

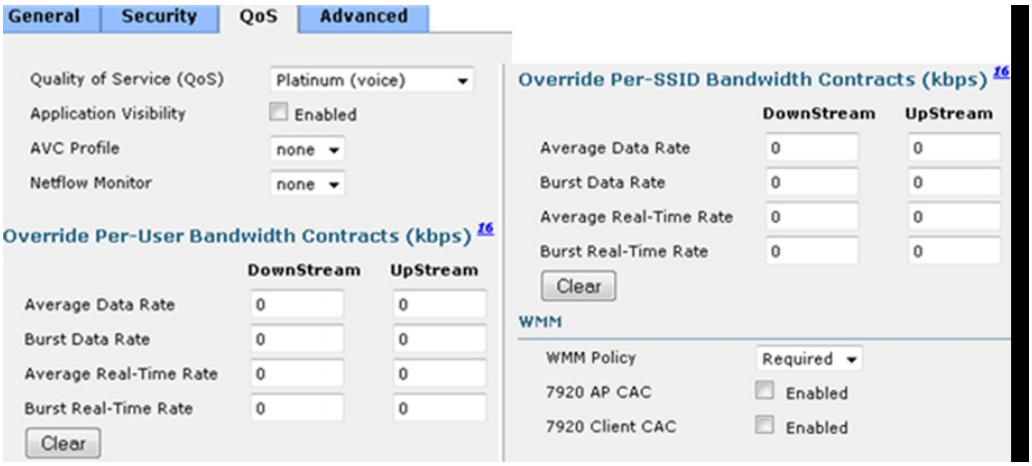

Each of the profiles (platinum, gold, silver, or bronze) is annotated with its typical use. In addition, a client can be assigned a QoS profile based on its identity, through authentication, authorization, and accounting (AAA). For a typical enterprise, the WLAN deployment parameters, such as per-user bandwidth contracts and over-the-air QoS, should use the default values, and standard QoS tools, such as WMM and wired QoS, must be used to provide optimum QoS to clients.

In addition to the QoS profiles, you can also control the WMM policy per WLAN as shown in the preceding figure with the following options:

- **Disabled:** WLAN does not advertise WMM capabilities or allow WMM negotiations
- **Allowed:** WLAN allows WMM and non-WMM clients
- **Required:** Only WMM-enabled clients are associated with this WLAN

QoS basic service set

The following figure shows the QoS basic service set (QBSS) information element (IE) that a Cisco AP recommends. The **Load** field indicates the portion of available bandwidth that is currently used to transport data on that AP.

Figure 35: QBSS Information Element

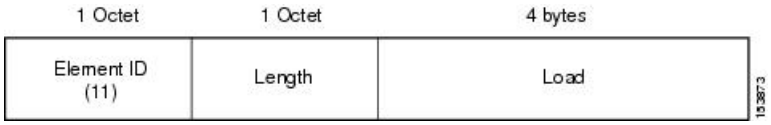

The QBSS in use depends on the WMM and clients settings on the WLAN. Based on the requirements, the following three types of QBSS IEs must be support:

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

Figure 34: WLAN Controller WLAN Default QoS Profile Settings, on page 66 shows 7920 AP and Client CAC, components of WLAN Controller (WLC) WLAN configuration that enables the AP to include appropriate QBSS elements in its beacons. WLAN clients with QoS requirements, such as the Cisco Unified Wireless IP Phones, use these recommended QoS parameters to determine the best AP with which to associate.

The WLC provides 7920 CAC support through the client call admission control (CAC) limit, or AP CAC limit. These features are listed below:

- **Client CAC limit:** The 7920 Client CAC maps to the old QBSS method, which is not clear channel assessment (CCA) based, but only accounts for 802.11 traffic on that specific AP. The client can set a fixed CAC limit, to prevent outbound calls when that limit is reached.
- **AP CAC limit:** The 7920 AP CAC maps to the new QBSS method, which is CCA-based, and accounts for all energy on the RF channel including 802.11 traffic for the local AP as well as for other APs, and also energy from non-802.11 devices (for example, microwaves and Bluetooth). The client can set a fixed CAC limit, to prevent outbound calls when that limit is reached.

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

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

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

Setting the admission control parameters

The following figure shows a sample configuration screen for setting the voice parameters on the controller.

Figure 36: Voice Parameter Setting

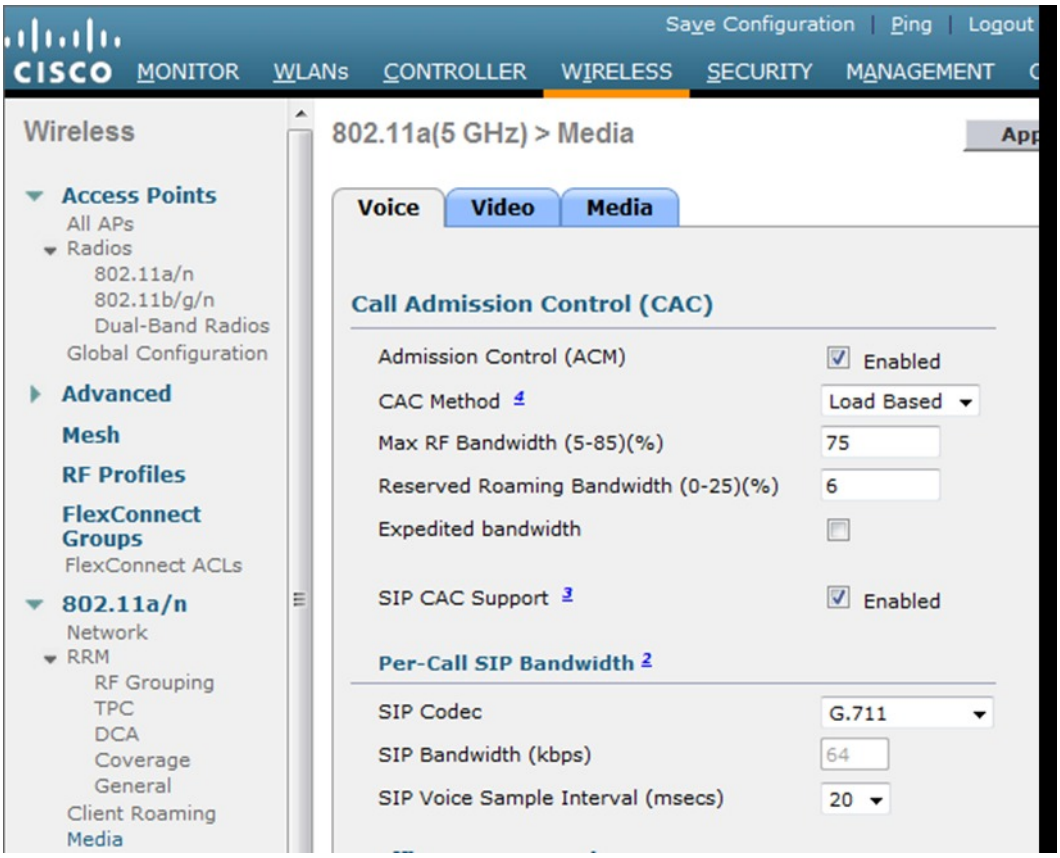

For SIP-based Cisco WLAN endpoint and mobile client deployments, Cisco recommends not to enable SIP CAC support because they utilize TCP-based SIP versus UDP-based SIP.

The admission control parameters consist of the maximum RF Bandwidth that a radio can use and still accept the initiation of a voice or video over WLAN call through a normal ADDTS request. The reserved roaming bandwidth is the capacity set aside to respond to the ADDTS requests during association or re-association, which are RToWLAN clients with calls in progress trying to roam to that AP.

Check the **Admission Control (ACM)** check box to enable admission control based on these parameters. This enables admission control based upon the AP capacity, but does not consider the possible channel loading impact of other APs in the area. To include this channel load in capacity calculations, select Load Based from the **CAC Method** drop down and check the **Admission Control (ACM)** check box.

The following figure shows an example of voice statistics reports that are available on the WCS, which displays the calls that are established on the radio of one AP, and the number of calls that roamed to that AP. This

report and other voice statistics can be scheduled or ad hoc, and either graphically displayed or posted as a data file.

Figure 37: Voice Statistics from WCS

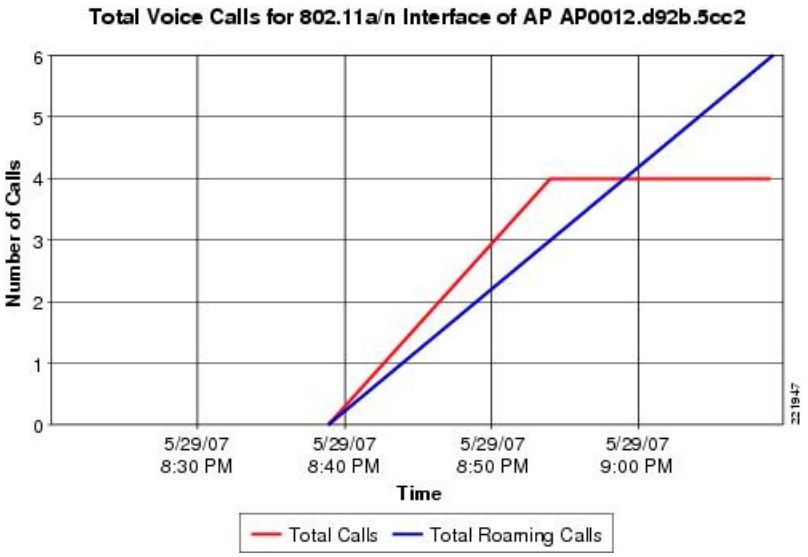

Note Call admission control is performed only for voice and video QoS profiles.

Impact of TSPEC admission control

The purpose of TSPEC admission control is to protect the high-priority resources. Therefore, a client that has not used TSPEC admission control does not have its traffic blocked but its traffic is reclassified if it tries to send, which it must not do if the client is transmitting WMM-compliant traffic in a protected AC.

The following tables describe the impact on classification if admission control is enabled, depending on whether a traffic stream has been established.

Table 8: Upstream Traffic

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

Table 9: Downstream Traffic

	Traffic Stream Established	No Traffic Stream
No admission control	No change.	No change.
Admission control	No change.	Packets are re-marked to BE (both CoS and DSCP) before they enter the network for WMM clients. For non-WMM clients, packets are sent with WLAN QoS.

IEEE 802.11e, IEEE 802.1P, and DSCP mapping

In a Unified Wireless network, WLAN data is tunneled through CAPWAP (IP UDP packets). To maintain QoS classification applied to WLAN frames, a process of mapping classifications to and from DSCP to CoS is required.

For example, when a WLAN client sends WMM classified traffic, it has an IEEE 802.1P classification in its frame. The AP must translate this classification into a DSCP value for the CAPWAP packet carrying the frame, to ensure the packet is treated with the appropriate priority as it reaches WLC. A similar process must occur on the WLAN Controller (WLC) for CAPWAP packets going to the AP.

A mechanism to classify traffic from non-WMM clients is also required, to ensure the AP and WLC give an appropriate DSCP classification to the CAPWAP packets for non-WMM clients.

The following figure shows a numbered example of the traffic classification flow for a WMM client, an AP, and a WLC.

Figure 38: WMM and IEEE 802.1P Relationship

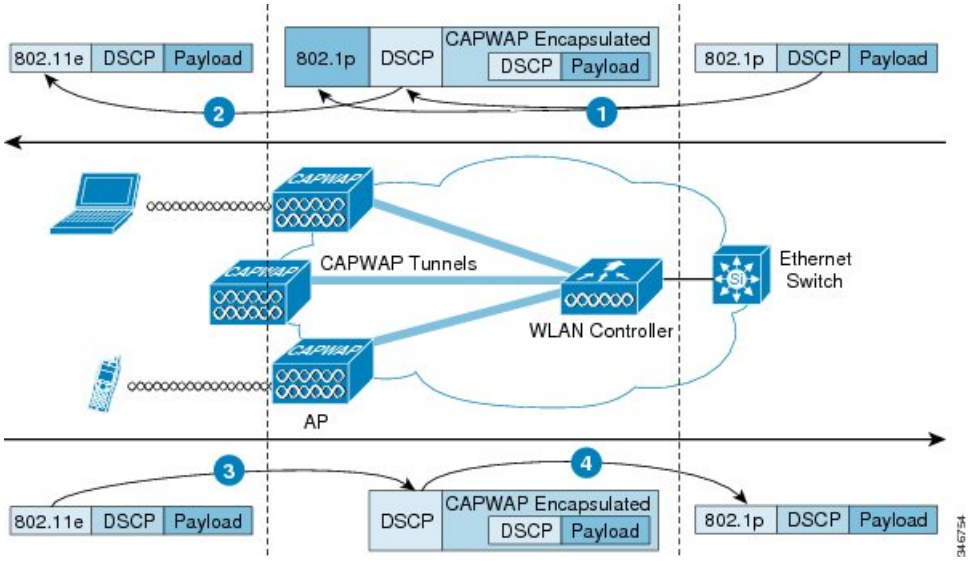

The traffic classification flow is described as follows:

1. A frame with an 802.1P marking and a packet with an IP DSCP marking arrive at the WLC wired interface. The IP DSCP of the packet determines the DSCP of the CAPWAP packet leaving the WLC.
2. The IP DSCP of the CAPWAP packet reaching the AP translates to an 802.11e CoS marking.
3. The 802.11e CoS marking of a frame arriving at the AP translates to an CAPWAP DSCP value, capped at the maximum value for that QoS profile.
4. The DSCP of the packet leaving the WLC will be equal to the DSCP of the packet that left the WLAN client. The 802.1P value of the frame depends on:
   - QoS translation table (see Table 10: Access Point QoS Translation Values, on page 72)
   - QoS profile for the WLAN
   - Wired QoS protocol that is configured for that QoS profile (see Figure 33: Editing WLAN Controller QoS Profiles, on page 65)

   If the wired QoS protocol is configured as *None*, then no 802.1p value is set. But if the protocol is set to 802.1p, then the 802.1p used depends on the translation table capped at a maximum value of 802.1p table value.

The multiple classification mechanisms and client capabilities require multiple strategies:

- CAPWAP control frames require prioritization, and CAPWAP control frames are marked with a DSCP classification of CS6.
- WMM-enabled clients have the classification of their frames mapped to a corresponding DSCP classification for CAPWAP packets to the WLC. This mapping follows the standard IEEE CoS-to-DSCP mapping, with the exception of the changes that are necessary for QoS baseline compliance. This DSCP value is translated at the WLC to a CoS value on IEEE 802.1Q frames leaving the WLC interfaces.
- Non-WMM clients have the DSCP of their CAPWAP tunnel set to match the default QoS profile for that WLAN. For example, the QoS profile for a WLAN supporting wireless IP phones would be set to Platinum, resulting in a DSCP classification of EF for data frames packets from that AP WLAN.
- CAPWAP data packets from the WLC have a DSCP classification that is determined by the DSCP of the wired data packets that are sent to the WLC. The AP table converting DSCP to WMM classification determines the IEEE 80211.e classification used when sending frames from the AP to a WMM client.

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

QoS baseline priority mapping

The CAPWAP AP and WLC perform QoS baseline conversion to ensure that WMM values are mapped to the appropriate QoS baseline DSCP values, rather than the IEEE values.

Table 10: Access Point QoS Translation Values

Traffic Type	IP DSCP	QoS Profile	802.1p	IEEE 802.11e UP
Inter-network control (CAPWAP control, 802.11 management)	48 (CS6)	Platinum	6	7
Voice	46 (EF)	Platinum	5	6
Interactive video	34 (AF41)	Gold	4	5
Mission critical	26 (AF31)	Gold	3	4
Call signaling	24 (CS3)	Gold	3	4
Transactional	18 (AF21)	Silver	2	3
Bulk data	10 (AF11)	Bronze	1	2
Best effort	0 (BE)	Silver	0	0
Scavenger	2	Bronze	0	1

The following table shows the translations values if AP is translating CoS values, for example, autonomous APs.

Table 11: WMM Packet Re-marking for APs with Priority Type Configured

Downstream L2 Packet Re-marking [1]			Upstream L2 Packet Re-marking		
Typical Application	CoS	WMM UP	802.1d Designation	WMM UP	CoS
Best Effort Data	0	0	BE	0	0
Medium Priority Data	1	2	BK	1	1
High Priority Data	2	3	-	2	1
Call Signaling	3	4	EE	3	2
Video Conferencing	4	5	CL	4	3
Voice Bearer	5	6	VI	5	4

Downstream L2 Packet Re-marking [1]			Upstream L2 Packet Re-marking		
Typical Application	**CoS**	**WMM UP**	**802.1d Designation**	**WMM UP**	**CoS**
Reserved	6	7	VO	6	5
Reserved	7	7	NC [2]	7	7

1 **In the downstream direction, the AP takes CoS markings on the wired interface and maps them to the UPs shown. In the upstream direction, the AP takes UPs that are received on the dot11 interface and maps them to CoS on the wired interface. Using this remapping results in the best match of WMM AC to CoS.**

2 The only network control traffic that must get mapped to CoS=7 is spanning-tree traffic that is used when work group bridges are deployed or when outdoor bridges are deployed, connecting the LANs between two or more buildings. Even though 802.11 MAC management traffic is carried on UP=7 in autonomous APs, is it not bridged onto the wired port of the AP.

Deploying QoS features on CAPWAP-based APs

Consider the following when you deploy QoS on wireless APs:

- The wired CAPWAP AP interface reads or writes Layer 2 CoS (IEEE 802.1P) information. The WLC and APs depend on Layer 3 classification (DSCP) information to communicate WLAN client traffic classification. The intermediate routers can modify this DSCP value and therefore the Layer 2 classification that is received by the destination does not reflect the Layer 2 classification that is marked by the source of the CAPWAP traffic.
- The APs no longer use NULL VLAN ID. As a result, L2 CAPWAP does not effectively support QoS because the AP does not send the IEEE 802.1P/Q tags, and in L2 CAPWAP there is no outer DSCP on which to fall back.
- APs do not reclassify frames; they prioritize based on CoS value or WLAN profile.
- APs use EDCF-like queuing on the radio egress port only.
- APs use first-in first-out (FIFO) queuing only on the Ethernet egress port.

WAN QoS and FlexConnect

For WLANs that have data traffic forwarded to the WLC, the behavior is same as non-FlexConnect APs (formerly hybrid remote edge access point or H-REAP) APs. For locally switched WLANs with WMM traffic, the AP marks the 802.1P value in the 802.1Q VLAN tag for upstream traffic. This occurs only on tagged VLANs; that is, not native VLANs.

For downstream traffic, FlexConnect uses the incoming 802.1Q tag from the Ethernet side and uses this to queue and mark the WMM values on the radio of the locally switched VLAN.

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

Wireless QoS deployment guidelines

The guidelines that you consider when you deploy QoS in a wired network apply when you deploy QoS in a wireless network. QoS does not create additional bandwidth; it prioritizes and optimizes the bandwidth that is allocated to different applications.

Successful wireless QoS requires awareness of the types of traffic and protocols traversing the network, and understanding of the specific delay sensitivity and bandwidth requirements of applications to properly design and configure WLAN QoS.

Throughput

It is important to consider and understand the offered traffic when you deploy IEEE 802.11 QoS. You must consider both bit rate and frame size, because IEEE 802.11 throughput is sensitive to the frame size of the offered traffic.

The following table shows how frame size affects throughput; a decrease in the packet size decreases the throughput.

Table 12: Throughput compared to frame size

	300	600	900	1200	1500	Frame size (bytes)
11g/a 6-54 Mbps	11.4	19.2	24.6	28.4	31.4	Throughput Mbps
11b 1-11 Mbps	2.2	3.6	4.7	5.4	6	Throughput Mbps

For example, if an application that offers traffic at a rate of 3 Mbps is deployed on an 11 Mbps IEEE 802.11b network, but uses an average frame size of 300 bytes, no QoS setting on the AP allows the application to achieve its throughput requirements. This is because IEEE 802.11b cannot support the required throughput for that throughput and frame size combination. The same amount of offered traffic, having a frame size of 1500 bytes, provides better throughput.

QoS switch configuration

This section discusses wired switch port configuration for the wired to wireless boundary of the following wireless infrastructure components:

- AP wired switch attachment
- WLC wired switch attachment

AP wired switch attachment

The QoS configuration of the AP switch is relatively trivial because the switch must trust the DSCP of the CAPWAP packets that are passed to it from the AP. There is no class of service (CoS) marking on the CAPWAP frames that come from the AP.

The use of IOS command **mls qos trust dscp** at the access switch enables trust of DSCP markings of the AP as set by the WLC policy. The maximum DSCP value that is assigned to client traffic is based on the QoS policy that is applied to the WLANs on that AP.

The above configuration command addresses only packet classification. Depending on local QoS policy, you can add queuing commands and other QoS-related configuration.

WLC wired switch attachment

The QoS classification on a WLC-connected switch is more complex than on the AP-connected switch because you must decide to trust either the DSCP or the CoS of traffic coming from the WLC. The following factors help to decide on the QoS switch configuration:

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

The use of the **mls qos trust cos** IOS command enables the trust of the CoS settings of the WLC. This allows a central location for the management of WLAN QoS, rather than managing the WLC configuration and an additional policy at the WLC switch connection. Customers who want a more precise degree of control can implement QoS classification policies on the WLAN-client VLANs.

Application visibility and control (AVC) for wireless

Cisco wireless AVC benefits include:

- Improved quality of experience for all wireless users through application-level optimization and control.
- Proactive monitoring and end-to-end application visibility to accelerate troubleshooting and minimize network downtime.
- Network capacity management and planning through greater visibility of application usage and performance.
- Prioritization of business-critical applications and subflows like Cisco Jabber voice or IM sessions.

AVC functionality

AVC on the Wireless LAN Controller has the functionality and features that are comparable to AVC found on other Cisco products. AVC application recognition is configurable down to the WLAN/SSID. Each WLAN can optionally enable various AVC parameters making any WLAN unique. WLANs define the name of the SSID with which the clients authenticate and associate. The same WLAN configuration also defines the highest level of Wi-Fi QoS for packet transmission over the Wi-Fi channel. A packet remarked by an AVC profile will not have a QoS priority that is above the QoS priority that is defined by the WLAN setting. For example, if WLAN is created for guest users with the QoS priority level of best-effort. Voice and video packets are transmitted at a best-effort priority, even if the AVC policy recognizes the packet as an audio packet. The guest SSID can be configured to limit “FaceTime” calls to the best-effort priority. Also, the same guest SSID

can be configured to use an AVC profile that blocks “You Tube” thus providing more bandwidth on other SSIDs that share the same Wi-Fi channel.

For Wi-Fi clients that are associated to a WLAN, AVC on the Wireless LAN Controller uses application recognition through deep packet inspection to determine how packets of a particular application should be handled per the Wireless LAN Controller AVC configurations. The Wireless LAN Controller is the control point for blocking packets and changing the QoS marking of packets. You can block the FaceTime application from connecting to the servers that establishes a FaceTime call. The blocking occurs at the Wireless LAN Controller. When the AVC profile blocks an application, the client device remains associated to the WLAN. If AVC profile is created to remark the FaceTime application packet, the remarking occurs at the Wireless LAN Controller. The remarking is done in the upstream and downstream direction. In the case of upstream traffic (from Wi-Fi endpoint to Wireless LAN Controller through the AP), the packet remarking is from the Wireless LAN Controller to the packets that are being forwarded to the destination endpoint. AVC cannot control the QoS packet markings at the source client or the markings of those packets because they are forwarded from the AP to the Wireless LAN Controller. In the case of downstream traffic (endpoint packets being forwarded by the Wireless LAN Controller through the AP to the Wi-Fi endpoint), the packet remarking occurs at the Wireless LAN Controller. The AP forwards the FaceTime traffic to the WLAN with 802.11e/WMM QoS priorities that is representative of the DSCP values assigned in the AVC FaceTime profile.

CAPWAP is the protocol that connects the APs and the Wireless LAN Controllers. CAPWAP packets encapsulate IP application packets. CAPWAP QoS packet markings upstream are based on 802.11e/WMM QoS values of the Wi-Fi header of the endpoint application packet. CAPWAP QoS packet markings downstream are based on the WLAN configurations. In the FaceTime example, the DSCP values on the header of the CAPWAP packets are assigned at the Wireless LAN Controller by the DSCP values that are configured in the AVC profile for FaceTime. You can configure AVC profiles for each Wireless Lan Controller and assign to the WLANs.

Note AVC configuration options for Wireless LAN Controller Version 7.4 and higher are provided in the Wireless LAN Controller configuration guides by Wireless LAN Controller release code version numbers. You can download the Wireless LAN Controller release configuration guides from Cisco.com. A separate WLC/AVC configuration guide is available by Wireless LAN Controller hardware type on the same Cisco product page as the Wireless LAN Controller software page.

AVC versions

The Wireless LAN Controller version of AVC runs as part of Wireless LAN Controller and does not require a separate license. AVC on the Wireless LAN Controller became available with Cisco Wireless LAN Controller Release 7.4.

Supported by the Wireless LAN Controller AVC are FTP/TFTP loads of Network Based Application Recognition (NBAR) protocol packs that are release matched to the NBAR engine version incorporated in the Wireless LAN Controller release. For example, the Wireless LAN Controller Release 7.5 uses NBAR engine version 13. Hence, protocol packs that are released for Release 7.5 will have a numbering that is similar to pp-AIR-7.5-13-4.1.1.pack.

You can determine the version of the protocol pack and AVC engine by executing the following Wireless LAN Controller CLI commands:

- **show avc protocol-pack version**
- **show avc engine version**

You can download the AVC NBAR2 Protocol Packs by the Wireless LAN Controller type from the same download page location as the software release versions for the Wireless LAN Controller posted on Cisco.com.

Traffic shaping, over-the-air QoS, and WMM clients

Traffic shaping and over-the-air QoS are useful tools in the absence of WLAN WMM features, but they do not help to prioritize IEEE 802.11 traffic directly. For WLANs that support WMM clients or wireless handsets, you must use the WLAN QoS mechanisms of these clients without using traffic shaping or over-the-air QoS.

Related Topics

- Enterprise QoS Solution Reference Network Design Guide
- Cisco Solutions for Enterprise Medianet: Optimizing Networks for Video, Voice, and Data
- Data Values for DiffServ Code Point and Type of Service Parameters
- Cisco AVC technologies and products

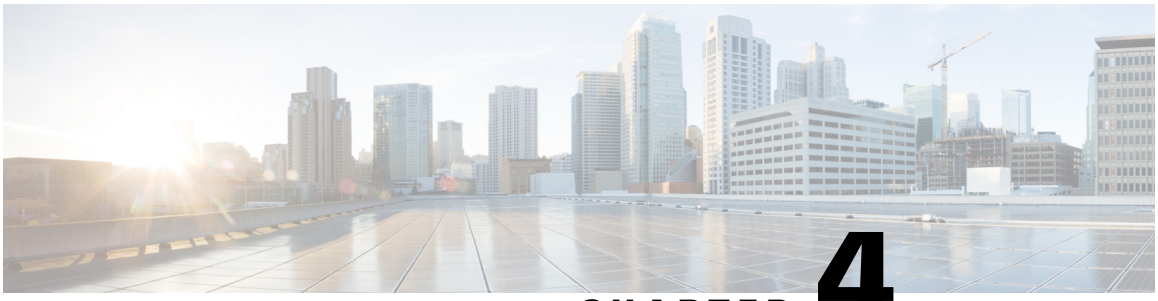

CHAPTER 4

Real-time Traffic over WLAN Security

The security of a wireless LAN (WLAN) system is always a critical consideration in every WLAN deployment. Control of the WLAN access depends on the principles of authentication, authorization, and accounting (AAA), augmented by encryption to ensure privacy. This chapter focuses on the authentication and encryption aspects of WLAN security for RToWLAN deployments.

For more information about WLAN security, see the *Enterprise Mobility Design Guide* at http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/emob73dg/emob73.html.

Real-Time Traffic over WLAN security overview

WLAN traffic is visible to any WLAN device within the radio frequency (RF) range of the WLAN infrastructure because the 802.11 wireless LANs is a shared network access medium.

The shared network access of WLANs creates the following challenges:

- How to provide privacy for users and devices of the WLAN from unauthorized users or devices.
- How to provide privacy for authorized users and devices of the WLAN from each other.
- How to provide privacy for multicast and broadcast WLAN traffic.
- How to differentiate between user and devices on the WLAN.

Each generation of WLAN security has addressed these challenges in different ways. But the key mechanisms are based on the same strategies that are used to secure communication over any untrusted medium—authentication and encryption.

802.11 security schemes

There are several 802.11 security schemes that you can implement when you deploy a real-time traffic enabled WLAN. The type of security scheme or schemes that the network administrator uses depends on the capabilities and features that are supported by the WLAN network infrastructure and the specific RToWLAN client devices that will be deployed.

The following figure shows the basic components of WLAN security. The components that are required to implement secure network attachment and encrypted traffic for wireless network traversal include the following:

- Wireless client device
- Wireless access point (AP)
- Wireless LAN Controller (WLC)
- Authentication or AAA server

Figure 39: Secure Wireless LAN Topology

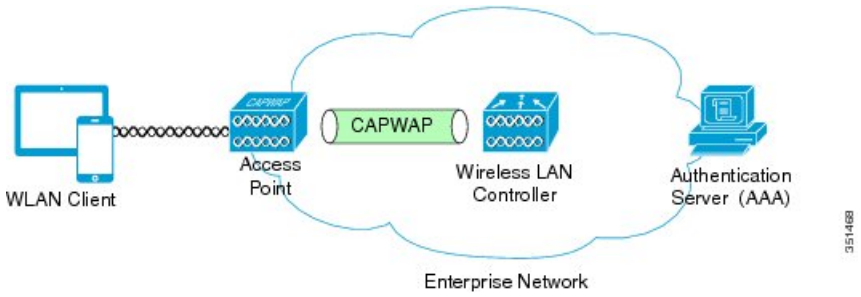

The best practice to deploy RToWLAN-capable client devices and RToWLAN services is that the security mechanisms that are enabled on the WLAN should provide user and device authentication and traffic encryption to ensure the following:

- Only authorized users and their devices are given access to the network.
- Real-time traffic flows are protected from interception and eavesdropping.

Open security scheme

An open security scheme provides no encryption or authentication for client device access to the WLAN.

The following figure shows an open WLAN security topology.

Figure 40: Open Wireless LAN Security Topology

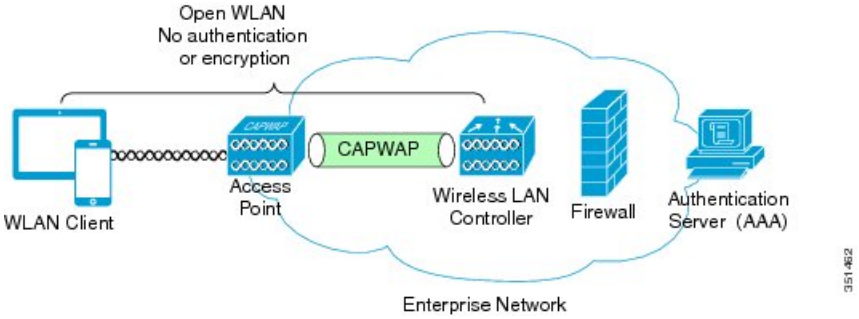

WLAN is open for attachment by all 802.11-capable devices. This security scheme is generally considered undesirable when deploying an enterprise WLAN, because the network is not protected from unauthorized access, and user and client traffic is not protected from interception and eavesdropping. However, despite the lack of encryption and authentication, open WLAN service set identifiers (SSIDs) are useful in some deployments when providing guest access for basic Internet connectivity or to onboard personal or noncorporate devices in bring-your-own-device (BYOD) scenarios.

In the case of BYOD deployments, the open network provides initial WLAN access for all clients and devices prior to identifying and onboarding authorized users and devices through additional network management and security services such that they have access to secured WLAN SSIDs or other areas of the enterprise network. You must take care to segment these open WLAN SSID networks from the rest of the secured enterprise network to prevent unauthorized access when you implement an open security scheme in these scenarios.

Wired Equivalent Privacy

A Wired Equivalent Privacy (WEP) security scheme provides encryption through a common shared key with minimal, if any, user or device authentication for client device access to the WLAN. The following figure shows a WEP security topology where WEP encryption occurs between the WLAN client device and the WLAN infrastructure AP and WLC.

Figure 41: WEP Wireless LAN Security Topology

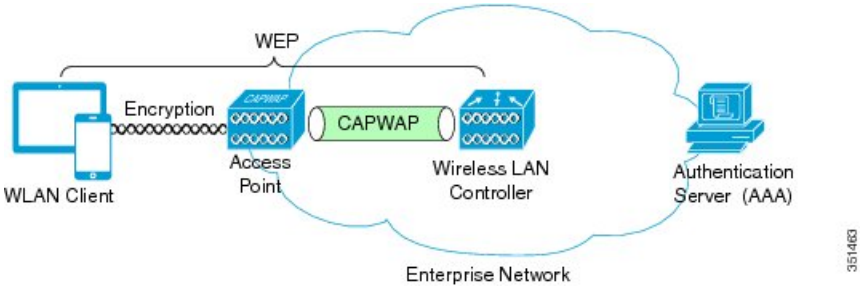

The original 802.11 standard defined the WEP encryption mechanism but did not define an authentication mechanism. The level of authentication that was offered in the original 802.11 standard was at a group level that required everyone in the group to have the same static encryption key. This key was used to encrypt unicast, multicast, and broadcast traffic. WLAN security solutions further augmented this group authentication

by authenticating the client MAC address. However, the solution of authenticating the client MAC address is not considered a significant improvement in security for the following reasons:

- It does not provide any additional per-user privacy, because the WEP key is still shared by all users.
- It is difficult to manage the WEP keys, because if one or multiple WEP keys need to be changed, you must update all the devices.
- It offers a weak level of authentication, because the 802.11 MAC addresses are sent unencrypted and the MAC address identifies the WLAN client devices rather than the users.
- It is difficult to administer MAC address authentication for large groups of users because you must maintain the database of the client device MAC addresses.

The original WEP encryption method with or without MAC address authentication that was implemented with the 802.11 standard were based on static configuration. While, the introduction of a dynamic WEP mechanism was an improvement over static WEP key implementations, issues in the WEP encryption mechanism means that the security of both static and dynamic WEP may be compromised. The issues of the WEP encryption mechanism itself is based on the fact that the WEP key can be derived by monitoring the client traffic.

Wi-Fi Protected Access

The weaknesses in WEP and the demand for a solution drove the Wi-Fi Alliance to develop WLAN security improvements through the 802.11i workgroup. These improvements are defined as Wi-Fi Protected Access (WPA). WPA addressed the main weakness in WEP encryption by replacing it with the Temporal Key Integrity Protocol (TKIP). While there are slight differences between WPA and the related sections of the 802.11i standard, these differences are transparent to users. The sections from the 802.11i standard that are used by WPA primarily address the need to secure the WLAN while maintaining sufficient backward compatibility with WEP to prevent the need to upgrade or replace currently deployed hardware. Because TKIP reuses the core encryption engine of WEP (RC4), it allows WPA to be implemented in the majority of systems through a software or firmware upgrade. WPA also attempts to address the absence of an authentication mechanism which was missing in the original WEP definition and further provides improvement over MAC address authentication that is sometimes used in combination with WEP.

Wi-Fi Protected Access 2

The security features that were developed in WPA were based on the recommendations of the 802.11i workgroup that was tasked with replacing original security features that were defined in the 802.11 standard. While the security changes from 802.11i that were adopted by WPA are important, the key component in 802.11i standard was the incorporation of the Advanced Encryption Standard (AES) into WLAN security that aligns its encryption mechanism with the latest industry standard for encryption. The underlying mechanism of AES-Counter Mode CBC-MAC (AES-Counter Mode describes the encryption mechanism, and CBC-MAC describes frame protection mechanism) is very different to those of WPA and WEP, and generally requires hardware upgrades to be supported. WPA2 introduces support for AES. The hardware requirements that are required to support AES encryption in WPA2 mean that migration from WPA is almost always dependent on a hardware upgrade or replacement. In many cases, it is easier to update the network infrastructure than to update the WLAN client infrastructure, and a complete migration to WPA2 is dependent on a generational change in the WLAN client infrastructure. Also, when you consider whether to migrate from WPA to WPA2, keep in mind that currently there are no known serious security exposures in WPA.

WPA and WPA2: Enterprise and Personal

The following figure shows the two different security schemes that the WPA and WPA2 define, WPA/WPA2-Personal and WPA/WPA2-Enterprise.

Figure 42: WPA /WPA2 Wireless LAN Security Topology

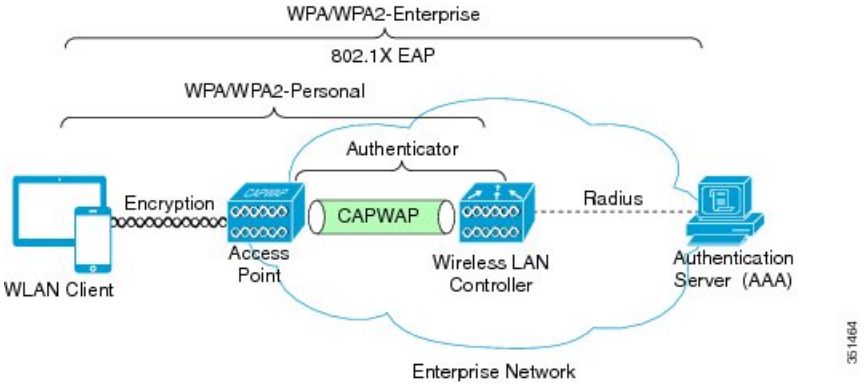

- **WPA/WPA2-Personal:**

  WPA/WPA2-Personal uses the same cryptographic tools as WPA/WPA2-Enterprise but relies on a shared key to authenticate the WLAN clients. The shared key mechanism of authentication that is used in WPA-Personal does not provide a per-user or per-device authentication. Therefore, every device and every AP that is part of that WLAN SSID uses the same shared key. On the other hand, the key that is used for encryption is unique per user and per session because of randomizing during the initial four-way cryptographic handshake, but the shared key that is used to authenticate is the same for everyone. The primary advantage of WPA/WPA2-Personal in an RToWLAN deployment is that it does not require the use of an AAA server, and this can be an advantage in smaller deployments or multisite deployments with one or more branch sites that are separate from a central or larger regional site.

  Note When you rely on WPA/WPA2-Personal security scheme, make sure that you use strong keys, because tools are available that can successfully perform a dictionary attack on WPA/WPA2-Personal.

- **WPA/WPA2-Enterprise:**

  WPA/WPA2-Enterprise uses the same base WPA frame protection and cryptographic features as WPA/WPA2-Personal, but adds 802.1X with Extensible Authentication Protocol (EAP)-based authentication to the scheme. 802.1X with EAP-based authentication requires utilization of an AAA authentication server.

WPA/WPA2-Enterprise versus WPA/WPA2-Personal

Generally, the use of WPA/WPA2-Enterprise is preferred over WPA/WPA2-Personal for enterprise RToWLAN deployments. WPA/WPA2-Personal is generally targeted more for home-user or small-office deployments. Shared key security systems do not provide the authentication features that are typically required for the enterprise, and can introduce operational issues due to the overhead in updating the shared keys if an RToWLAN client device is lost, stolen, or is part of a regular key rotation regime.

The reward for successfully cracking, guessing, or stealing the shared key is very high, because this key is used for all users and devices. This does not mean that you can not use WPA/WPA2-Personal for RToWLAN deployments. You must balance the enterprise security requirement for an AAA authentication server against the RToWLAN handset requirements and characteristics of the RToWLAN deployment. Real-time traffic services and applications are usually expected to be highly available, and high availability may be difficult to achieve in branch and remote environments that are dependent on a centralized authentication system. You can address this issue by distributing authentication databases to branches through local AAA authentication servers or the embedded AAA services of a Wireless LAN Controller; or by deploying RToWLAN on a WLAN system leveraging WPA/WPA2-Personal, given the lack of dependency on centralized authentication.

Related Topics

802.1X and Extensible Authentication Protocol, on page 84
802.11 encryption, on page 87

802.1X and Extensible Authentication Protocol

In order to provide enterprise-level WLAN security, 802.1X and Extensible Authentication Protocol (EAP) authentication mechanisms were implemented to provide mutual authentication of WLANs and WLAN client devices. The following figure shows the basic 802.1X and EAP authentication secure topology.

Figure 43: 802.1X and EAP Wireless LAN Security Topology

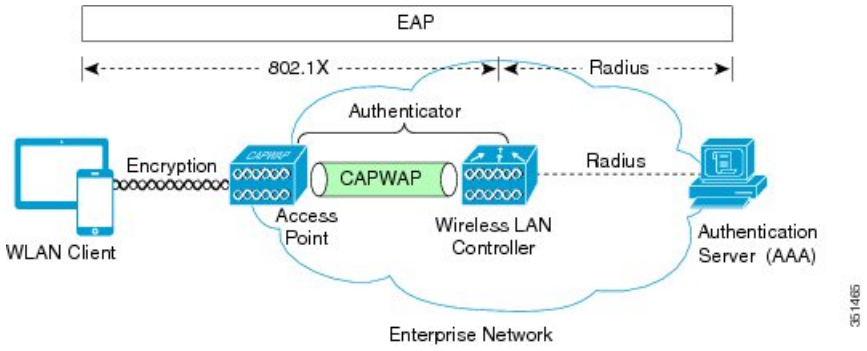

802.1X is an IEEE standard for port-based network access control and was adopted by the 802.11i security standard workgroup. The 802.1X standard provides authenticated access to 802.11 wireless LAN networks using the following logic:

- A virtual port is created at the AP during the 802.11 association process for each WLAN client device.
- The AP then blocks all data frames on this virtual port except for 802.1X-based traffic.
- EAP authentication packets are carried in 802.1X traffic frames and are passed by the AP and Wireless LAN Controller to the AAA authentication server.
- Assuming EAP authentication is successful, the authentication server sends an EAP success message back to the Wireless LAN Controller and AP, which in turn pass the message on to the WLAN client device.
- The AP then allows data traffic (including voice and video) from the WLAN client device to flow through the virtual port.

- Before the virtual port opens to allow data traffic, data link encryption is established between the client device and the AP.

During the authentication process, a unique per-user per-session shared key is derived, and a portion of this key is used as a per-session encryption key.

The EAP authentication process supports a number of protocols, and which protocol is used ultimately depends on the capabilities of the WLAN client device supplicant and the WLAN infrastructure. Regardless of the EAP type that is used, all protocols generally behave as shown in the example EAP flow in the following figure.

Figure 44: EAP Protocol Flow

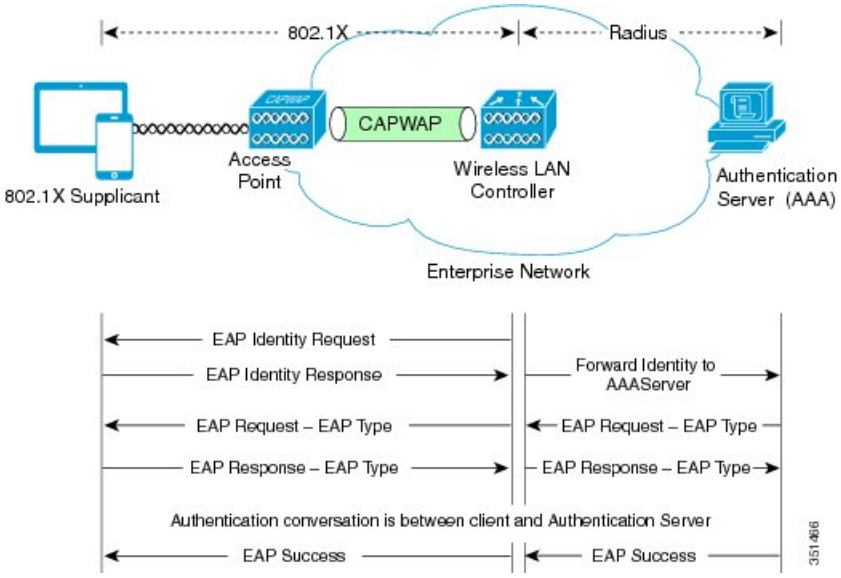

EAP as defined by RFC 3748 supports four packet types as part of the EAP authentication process:

- **EAP request:**

  The request packet that is sent by the authenticator (in the preceding figure, it is the Wireless LAN Controller and AP in combination) to the 802.1X supplicant (in the preceding figure, it is the WLAN client device). Each EAP request has a specific type that indicates what is being requested. In the example in the preceding figure, the first EAP request is for the WLAN client device identity, while the second EAP request is for the EAP type to be used for the authentication. A sequence number allows the authenticator and the peer to match an EAP response to each EAP request.

- **EAP response:**

  The response packet is sent by the WLAN client device to the AP and in turn to the Wireless LAN Controller, and uses a sequence number to match the initiating EAP request. In the case of an identity or type response, the response is forwarded by the Wireless LAN Controller to the authentication server.

- **EAP success:**

  Assuming that the WLAN client device or user has provided appropriate credentials during the authentication conversation, as shown in the preceding figure, the AAA server sends an EAP success packet to the Wireless LAN Controller, which in turn relays it through the AP to the WLAN client device.

- **EAP failure:**

  If the appropriate credentials are not provided at the WLAN client device or some other failure occurs, the AAA server sends an EAP failure packet to the Wireless LAN Controller, which relays it through the AP to the WLAN client device, resulting in failure of the authentication.

Common RToWLAN EAP supplicant types

This section describes the following common RToWLAN EAP supplicant types:

- EAP-FAST
- EAP-TLS
- PEAP

EAP-FAST

The EAP-Flexible Authentication via Secure Tunneling (EAP-FAST) protocol was meant as a replacement for pre-802.11i Cisco proprietary Lightweight EAP (LEAP). LEAP was specifically designed to consider the limited processing power of application-specific or purpose-built devices, such as RToWLAN IP handsets. Given the security issues of LEAP, where weak passwords (less than ten characters) could be derived through analysis of the LEAP authentication transactions, EAP-FAST was designed to address these security issues while at the same time maintaining the "lightweight" nature of LEAP.

EAP-FAST is designed to provide the same tunneling protection as a tunneled authentication protocol such as EAP-Transport Layer Security (TLS), without requiring the Public Key Infrastructure (PKI) overhead associated with setting up the TLS tunnel that is used in EAP-TLS. Instead EAP-FAST typically relies on protected access credentials (PACs) for authenticating the tunnel between the client device and the authentication server. Although automatic PAC provisioning may be used, if the PAC is intercepted, it can be used to access user credentials. The use of manual PAC provisioning or optionally authentication server certificates during provisioning can help mitigate this potential issue.

As a tunneled protocol, EAP-FAST is capable of supporting multiple inner authentication mechanisms such as Microsoft Challenge-Handshake Authentication Protocol Version 2 (MSCHAPv2) or generic token card (GTC). The supported inner authentication mechanism depends on the RToWLAN client implementation.

EAP-TLS

The EAP-TLS protocol provides tunneled authentication protection relying on the PKI to authenticate both the WLAN client device and the WLAN network infrastructure. EAP-TLS uses certificates for both user and server authentication and for dynamic session key generation. It requires installation of both a client certificate and an authentication server certificate. EAP-TLS provides excellent security but requires client certificate management.

PEAP

Protected EAP (PEAP) uses TLS to protect authentication exchange between the WLAN client device and the authentication server. In the case of PEAP MSCHAPv2, MSCHAPv2 is used to encapsulate this authentication exchange across the tunnel. With PEAP GTC, a generic token card exchange protects the authentication process across the tunnel.

802.11 encryption

Encryption is a critical component of 802.11 WLAN security and is necessary for providing privacy over a local RF broadcast network. RToWLAN deployments should leverage TKIP or AES encryption as part of the WPA or WPA2 security mechanisms for securing network communication whenever WLAN client devices and infrastructure permit because of the superior security that these advanced mechanisms provide. WPA and WPA2 security mechanisms provide significant improvements over WEP encryption because of the encryption key derivation method.

The encryption key with WPA and WPA2 is derived using a four-way cryptographic handshake where the key shared between the WLAN client and the WLAN AP is not used directly for encryption, but instead used as the basis for the four-way handshake which derives the encryption key. This four-way handshake is used in both WPA/WPA2-Personal and WPA/WPA2-Enterprise.

The keys that are used for encryption are derived from the pair-wise master key (PMK) that has been mutually derived during the EAP authentication. This PMK is sent to the authenticator in the EAP success message, but is not forwarded to the supplicant because the supplicant has derived its own copy of the PMK.

The following figure shows the basic four-way handshake mechanism that is used in WPA/WPA2-Enterprise.

Figure 45: Four-Way Handshake for Deriving Wireless Encryption Keys with WPA/WPA-2 Enterprise

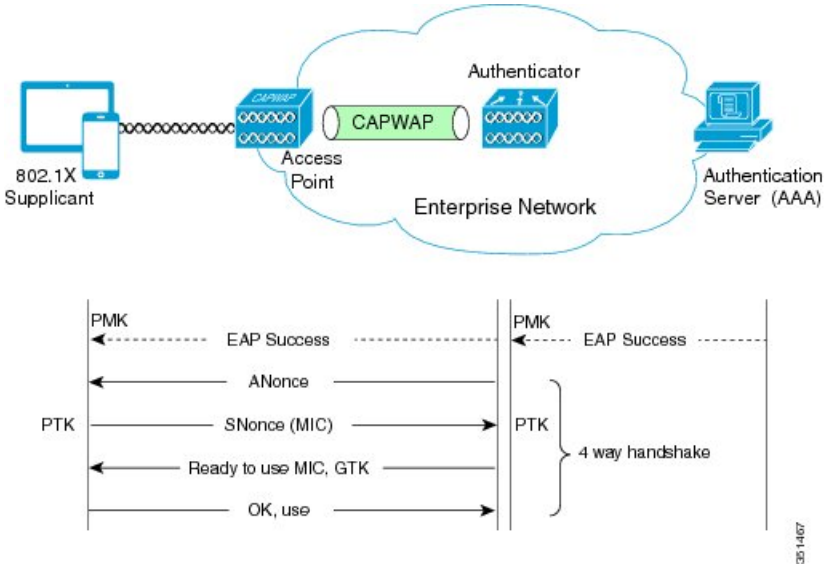

1. The authenticator sends an EAP over LAN (EAPOL)-Key frame that contains an authenticator nonce (ANonce), which is a random number that is generated by the authenticator.
   a. The supplicant generates a supplicant nonce (SNonce), which is a random number that is generated at the supplicant.
   b. The supplicant derives a pair-wise temporal key (PTK) from the ANonce and SNonce.
2. The supplicant sends an EAPOL-Key frame that contains the SNonce and a message integrity check (MIC).
3. The authenticator derives the PTK from the ANonce and SNonce and validates the MIC in the EAPOL-Key frame.

4 The authenticator sends an EAPOL-Key frame containing the group temporal key (GTK), the multicast, and the broadcast encryption keys if the validation is successful.

5 Upon validating the MIC from this frame, the supplicant installs its PTK and the GTK.

6 The supplicant sends an EAPOL-Key frame to confirm that the temporal keys are installed.

7 Upon validating the MIC from this frame, the authenticator installs the PTK for this client.

At this point, the supplicant and authenticator have verified that they both have a matching PMK, and both share the same PTK and GTK.

As shown in the preceding figure, with WPA/WPA2-Enterprise, the shared key that is used to generate the cryptographic key through the four-way handshake is derived during the 802.1X EAP authentication process. This EAP authentication process provides the AAA features that are missing in WPA/WPA2-Personal, allowing each user or device to be authenticated individually, a policy based on the authentication ID applied (authorization), and the collection of statistics based on authentication ID (accounting).

The difference between WPA/WPA2-Enterprise and the WPA/WPA2-Personal behavior is that the four-way handshake uses a shared key that is configured in the WLAN client device supplicant and the WLC. This shared key mechanism of authentication that is used in WPA-Personal does not provide a per-user or per-device authentication. Therefore, every device and every AP that is part of that WLAN SSID uses the same shared key. On the other hand, the key that is used for encryption, just as with WPA/WPA2-Enterprise, is unique per-user and per-session because of randomizing during the initial four-way cryptographic handshake.

Key caching and management

After an RToWLAN client device successfully authenticates with the WLAN network and establishes an encryption mechanism for traffic flows, it is able to securely send and receive traffic through the AP that it is associated with. However, what happens when the RToWLAN device moves or roams from one part of the WLAN network to another and must associate to another AP on the WLAN? In these situations, if the RToWLAN client device user is on an active voice or video call, it is important that the association of the client device to the new AP occurs as quickly as possible while at the same time maintaining implemented security authentication and encryption mechanisms. To facilitate this rapid reassociation and reauthentication, authentication and encryption keys must be managed and in some cases cached.

Related Topics

Real-time Traffic over WLAN Roaming, on page 91

Additional 802.11 security mechanisms

RToWLAN network administrators should consider additional 802.11 WLAN security mechanisms to prevent unauthorized access or disruptive network attacks apart from secure client device WLAN association, authentication, and traffic encryption. For example, traditional wired network attack vectors such as MAC flooding, man-in-the-middle attacks, and DHCP snooping or starving should be mitigated with appropriate network management and wired and wireless LAN infrastructure security features. Likewise, wired and wireless intrusion prevention, detection, and mitigation are critical components for a successful RToWLAN deployment. In particular, the detection and mitigation or elimination of rogue access points and clients is critical to maintain healthy wireless LAN radio frequency deployments with minimal interference. Without rogue AP and client detection, proper radio frequency design can be compromised, resulting in poor wireless

network throughput and capacity, unacceptable voice and video quality, and in some cases complete failure of real-time traffic applications and services.

Related Topics

Enterprise Mobility Design Guide

RToWLAN design considerations

An important aspect of the overall network design is to secure the wireless network for RToWLAN deployments. You must consider authentication and encryption method selection, scalability, and high availability aspects when you implement RToWLAN applications and services.

Authentication and encryption method selection

It is critical that you enable appropriate 802.11 security mechanisms for RToWLAN applications and services for a successful deployment. The most important considerations when you design a secure RToWLAN deployment are the capabilities of the WLAN infrastructure components such as access points, wireless LAN controllers, and the WLAN client devices themselves. Consider the following important factors:

- Always try to implement the strongest security mechanisms. WPA or WPA2-Enterprise security methods with 802.1X EAP authentication and TKIP or AES encryption is preferred over WEP or open authentication and encryption methods, provided the planned client devices and infrastructure are capable of supporting the stronger security methods. On the other hand, if the infrastructure or the target client devices do not support more secure mechanisms, then the network administrators must determine the most secure mechanism that can be supported based on the security policies and equipment standards of the organization.
- In the case of BYOD implementations where noncorporate-owned devices may be utilizing the WLAN infrastructure, it may be necessary to enable open authentication and nonencrypted connections to onboard these devices or simply to provide Internet-only guest access. In these cases, unsecure network access (at least initially) may be as critical as highly secure authentication and encryption methods.
- Another important consideration when selecting the security method or methods for an RToWLAN deployment is whether network attachment requires end-user intervention to complete the connection. The network administrator should strongly consider allowing voice and video-capable clients to attach to the enterprise network in the background (after initial provisioning), without user intervention. This ensures maximum utilization of the real-time traffic applications and services such as the enterprise voice and video telephony infrastructure. Specifically, the use of a certificate-based identity and authentication security mechanism like EAP-TLS helps facilitate an excellent user experience by minimizing network connection and authentication delay and ensuring quick WLAN device attachment. Failure to provide a fast, secure, and seamless network attachment mechanism for RToWLAN client devices may result in limited use of real-time traffic applications and services, because users delay or forget to intervene to complete the authentication process.
- Ultimately, the encryption and authentication methods that are used will be dictated by the client devices and the wireless infrastructure that will be deployed as well as the intended use cases for the RToWLAN deployment.

Scalability

While administrators generally prefer stronger security mechanisms, they must also consider the scalability of the security solution in terms of the number of users and the devices that will be deployed. In scenarios

where large numbers of users or devices are deployed, it is important that the authentication servers or services are able to handle the authentication request load during the busiest times for wireless-client-device-network attachment.

For example, at the beginning of a work day, depending on the size of the deployment, the number of users who are attempting to connect their RToWLAN devices to the wireless network may exceed the authentication or credential database storage capacity of the authentication server. This results in failed or delayed authentication for at least some of the devices which can be problematic. In these situations, the administrators should seek to distribute authentication fulfillment and credential storage across multiple authentication servers. Ensure that sufficient capacity is available to handle the expected authentication load for successful RToWLAN deployments.

High availability

Another important aspect of securing RToWLAN deployments is to ensure that WLAN network security services are highly available. For example, in situations where an authentication server is required to secure the RToWLAN client device network attachment, it is important that the authentication server is available to authenticate the device when it attempts to connect. In the situation where the authentication server has failed or is not available due to a network issue, it is critical that a redundant authentication server is available to handle the authentication request.

Without highly available authentication services and server redundancy, RToWLAN client devices may not be able to connect to the WLAN network during a server failure. For example, in distributed network deployments with multiple sites where the authentication services for a branch location are provided through a centralized authentication server located in the central site, if there is a network failure between the branch location and the central site, the client devices that are located in the branch site will not be able to authenticate and connect to the WLAN unless a redundant authentication service is available at the local site or at another site that is reachable from the local site. Whenever possible, an RToWLAN deployment should be designed with highly available authentication services.

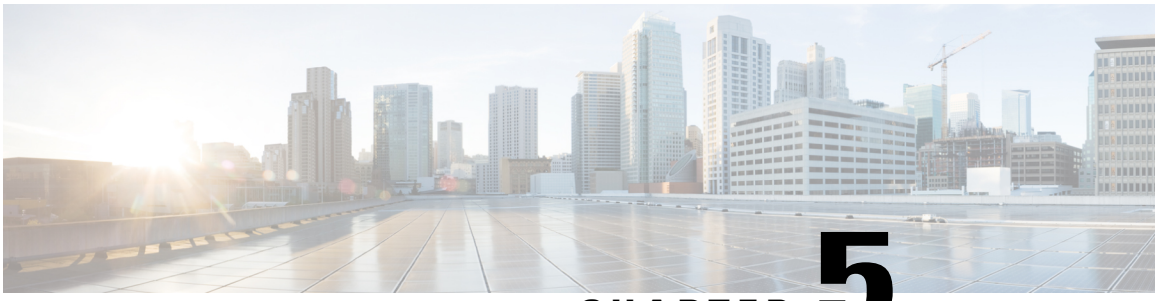

CHAPTER 5

Real-time Traffic over WLAN Roaming

At a basic level, *roaming* in an enterprise IEEE 802.11 network occurs when an IEEE 802.11 client changes its access point (AP) association from one AP to another AP within the same WLAN. Depending on client capabilities, an 802.11 WLAN client may roam on the same WLAN between APs within the same frequency band or between the 2.4 GHz and 5 GHz frequency bands. Smartphones and tablets that have simultaneous cellular and Wi-Fi connections may seamlessly roam across networks provided there is a suitable infrastructure network design. When a client roams from a WLAN with one service set identifier (SSID) to a WLAN with another SSID, the roam will not be seamless. The Wi-Fi client logic maintains only one Wi-Fi WLAN authentication at a time.

WLAN clients may roam based solely on their software capabilities or they may rely on assisted roaming capabilities provided by the WLAN infrastructure APs. In the case of client controlled roaming, the client is responsible to determine if it needs to roam, and then detects, evaluates, and roams to an alternative AP. The software that resides in the client evaluates the quality of the current Wi-Fi connection, and executes the connection and roam logic to join an alternate AP to gain a better quality connection.

Note WLAN standard bodies (such as the IEEE) and industry bodies (such as the Wi-Fi Alliance) do not specify when a client should roam, or how the client determines to which alternative AP it should roam. The roaming algorithms for each vendor are proprietary and are not generally published.

IEEE standards for 802.11r and 802.11k

Currently, IEEE 802.11k and 802.11r are the key industry standards for enabling seamless basic service set (BSS) transitions in the WLAN environment. The 802.11r and 802.11k standards support Wi-Fi 802.11r Fast Transition, secure authentication, and 802.11k neighbor list radio management.

With Cisco Unified WLAN controllers running release 7.4 or higher, mobile wireless devices running Apple iOS 6 and higher leverage 802.11k neighbor lists for enterprise roaming.

The following steps describe how an Apple iPhone requests, receives, and processes an 802.11k neighbor list:

1. The iPhone that is associated to an AP sends a request for a list of neighboring APs on the same WLAN. The request is in the form of an 802.11 management frame known as an action packet.
2. The AP responds with a list of neighboring APs on the same WLAN with their Wi-Fi channel numbers. This response frame is also an action packet.
3. The iPhone receives the response frame and identifies which APs are the entrants for upcoming roams.

The use of 802.11k radio resource management (RRM) process allows the mobile client device to roam efficiently and quickly. This is a requirement for good call quality in an enterprise environment where on-call roaming is common. As smartphone vendors adopt the 802.11r and 802.11k standards, their users can experience more efficient roaming along with good call quality during the roam.

The recommended WLAN controller (WLC) 802.11k configuration is to enable the RRM to provide both 2.4 GHz and 5 GHz AP channel numbers in the neighbor list response packets. Cisco recommends the use of 5 GHz band Wi-Fi channels for not only voice and video over WLAN calls but for all applications and devices.

With the neighbor list information, the mobile client device need not examine all of the 2.4 GHz and 5 GHz channels to find an AP it can roam to. This provides the following benefits:

- Reduces channel utilization on all channels, thus increasing bandwidth on all channels.
- Reduces roam times and improves the decision made by mobile devices.
- Increases battery life of the device because the device is neither changing the radio configuration for each channel nor sending probe requests on each channel.

The device does not have to process all of the probe response frames it receives on a channel. It only needs to validate that it can connect to an AP that is provided in the list of APs in the 802.11k neighbor list response frame.

Fast roaming

The recommended Enterprise security configuration for devices running Apple iOS 6 or higher is 802.11r Fast Transition. The IEEE 802.11r specification was approved in July 2008, and it follows the 802.11i specification of June 2004.

802.11r reduces the number of packets that are exchanged between the client and an AP. The client preauthenticates to the AP it will roam to before actually roaming. This means the roam itself occurs faster because the AP already has the client authentication credentials cached, resulting in fewer packets required between the client and the AP.

802.11r introduces the following standard-based fast transition:

- Allows a client to establish security and QoS state to roam-to AP before (or during) reassociation.
  - **Method 1: Over-the-Air (client to roam-to AP):** Exchanges four packets over the Wi-Fi channel.
  - **Method 2: Over the Distribution System (through the roam-from AP):** Exchanges two packets over the Wi-Fi channel and two packets through Ethernet

The following guidelines and limitations currently affect 802.11r Fast Transition:

- This feature is not supported on Mesh APs.
- For APs in FlexConnect mode:
  - 802.11r Fast Transition is supported only in centrally and locally switched WLANs in Cisco WLAN Release 7.3 and later.
  - This feature is not supported for the WLANs that are enabled for local authentication.
- This feature is not supported on Cisco 600 Series OfficeExtend Access Points.
- 802.11r client association is not supported on APs in standalone mode.
- 802.11r fast roaming is not supported on APs in standalone mode.
- 802.11r fast roaming is not supported between local authentication and central authentication WLANs.
- 802.11r fast roaming is not supported if the client uses over-the-distribution-system (DS) preauthentication in standalone mode. In over-the-DS roaming, packets are sent on the wired infrastructure.
- The service from a standalone AP to a client is only supported until the session timer expires.
- TSpec is not supported for 802.11r fast roaming.
- If a WLAN link latency exists, fast roaming is also delayed. The client must verify the voice or data maximum latency.
- The WLAN controller (WLC) handles 802.11r Fast Transition authentication requests during roaming for both over-the-air and over-the-DS methods.
  - Over-the-DS is recommended because two of the required packets are sent on the wired connection of the APs, with two packets sent on the WLAN. If you do not select the DS option, then all the four packets are sent on the WLAN.

Recommended WLAN controller configuration for fast transition

Use the following WLAN configuration recommendations to add 802.11r Fast Transition clients to the WLAN network.

Note These recommendations are the result of cooperative work between Apple and Cisco.

- Configure an additional WLAN for fast transition 802.1x clients.
- Configure an additional WLAN for fast transition PSK clients.
- Apple and Cisco recommend that you use separate WLAN and service set identifiers (SSIDs) for legacy clients.

The reason for these recommendations is that the legacy radio drivers cannot interpret the added information in the association response packets of a WLAN with fast transition configurations. Although the 802.11r specification was approved in the year 2008, not all client radio drivers have been updated to handle the changes in management packets with respect to 802.11r. This includes several Apple products.

Note The 802.11r specification changes the Wi-Fi packet structure. Legacy clients may not be programmed to accommodate the change and they fail to associate to a WLAN that enables 802.11r. Therefore it is recommended that you use a new WLAN for 802.11r-capable devices. iPad2 is an example of a device that cannot join an 802.11r WLAN.

The following figure shows a WLAN infrastructure with multiple WLANs and SSIDs to accommodate a range of client devices with varying specification support.

Figure 46: Example of Multiple WLANs and SSIDs

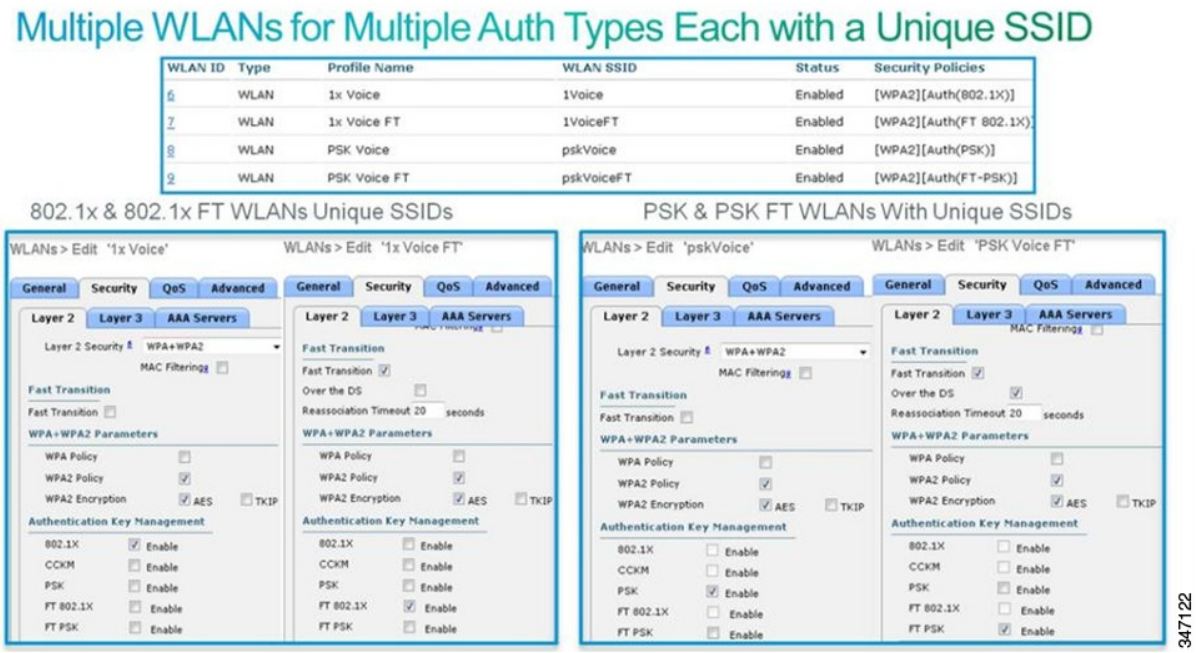

For information about command line interface (CLI) or graphical user interface (GUI) fast transition configuration options, see the *Cisco Wireless LAN Controller Configuration Guide* at http://www.cisco.com/ corresponding to the installed version of WLC firmware.

Related Topics

802.11k specification
Additional information about 802.11r
IEEE 802.11r specifications

Client roaming decision

802.11 wireless clients detect that roaming is required when the connection to the current AP degrades. Roaming necessarily affects client traffic because a client scans other 802.11 channels for alternative APs, reassociates, and authenticates to the roam-to AP. Before roaming, a client takes the following actions to improve its current connection without necessitating a roam:

- **Data retries:** The IEEE 802.11 MAC specifies a reliable transport. Every unicast frame that is sent between a wireless client and an AP is acknowledged at the MAC layer. The IEEE 802.11 standard specifies the protocol that is used to retry the transmission of data frames for which an acknowledgment was not received.
- **Data rate shifting:** IEEE 802.11a, 802.11b, and 802.11g each support a variety of possible data rates. The data rates that are supported for a given frequency band (for example, 2.4 GHz or 5 GHz) are configured on the wireless control system server (WCS) or WLC and are pushed down to the APs using that frequency band. Each AP in a given WLAN then promotes the supported data rates in its beacons. When a client or AP detects that a wireless connection is degrading, it can change to a lower supported transmission rate, because lower transmission rates generally provide superior transmission reliability.

Although the roaming algorithms differ for each vendor or driver version (and potentially for different device-types from a single vendor), the following common situations can typically cause a roam to occur:

- **Maximum data retry count is exceeded:** Excessive number of data retries are a common roam trigger.
- **Low received signal strength indicator (RSSI):** A client device can decide to roam when the receive signal strength drops below a threshold. This roam trigger does not require active client traffic to induce a roam.
- **Low signal-to-noise ratio (SNR):** A client device can decide to roam when the difference between the receive signal strength and the noise floor drops below a threshold. This roam trigger does not require active client traffic to induce a roam.
- **Proprietary load balancing schemes:** Some wireless implementations have schemes where clients roam in order to more evenly balance client traffic across multiple APs. This is one case where the roam may be triggered by a decision in the WLAN infrastructure and communicated to the client through vendor-specific protocols.

Cisco Compatible Extensions client roam triggers

Wireless LAN Controllers (WLC) are configured with a default set of RF roaming parameters that are used to set the RF thresholds that are adopted by the client to decide when to roam. You can override the default parameters by defining a custom set. These Cisco Compatible Extensions (CCX) parameters are defined on the WLC once for each IEEE 802.11 frequency band (2.4 GHz or 5 GHz).

WLAN clients that run on Cisco Compatible Extensions Version 4 or later are able to use the following parameters (which are communicated to the client through the Enhanced Neighbor List feature that is described in Cisco Compatible Extensions channel scanning, on page 97):

- **Scan threshold:** The minimum RSSI that is allowed before the client can roam to a better AP. When the RSSI drops below the specified value, the client must be able to roam to a better AP within the specified transition time. This parameter also provides a power-save method to minimize the time that the client spends in active or passive scanning. For example, the client can scan slowly when RSSI is above the threshold and scan more rapidly when RSSI is below the threshold.
- **Transition time:** The maximum time that is allowed for the client to detect a suitable neighboring AP to roam to and to complete the roam, whenever the RSSI from the clients associated AP is below the scan threshold. The scan threshold and transition time parameters guarantee a minimum level of client roaming performance. Together with the highest expected client speed and roaming hysteresis (for a definition of *hysteresis*,see below) these parameters help to design a WLAN network that supports roaming just by ensuring a certain minimum overlap distance between APs.

- **Minimum RSSI field:** A value for the minimum RSSI that is required for the client to associate to an AP.
- **Hysteresis:** A value to indicate how much greater the signal strength of a neighboring AP must be for the client to roam to that AP. This parameter is intended to reduce the amount of roaming between APs if the client is physically located on or near the border between two APs.
- **Call admission control (CAC):** A call admission control denial from the WLAN infrastructure can cause the client device to roam.

Note Even though a wireless client may be CCX compatible, it may still rely on 802.11k or its own proprietary roaming algorithm instead of the CCX triggers listed above.

Roaming selection of a new access point

Channel scanning

Wireless clients learn about available APs by scanning other 802.11 channels for available APs on the same WLAN or SSID. The wireless clients can scan other IEEE 802.11 channels in the following two ways:

- **Active scan:** Active scanning occurs when the client changes its 802.11 radio to the channel that is being scanned, broadcasts a probe request, and then waits to receive any probe responses (or periodic beacons) from APs on that channel (with a matching SSID). The 802.11 standards do not specify how long the client must wait, but 10 ms is a representative time period. The probe-request frames that are used in an active scan are of the following two types:
  - **Directed probe:** The client sends a probe request with a specific destination SSID; only APs with a matching SSID reply with a probe response.
  - **Broadcast probe:** The client sends a broadcast SSID (actually, a null SSID) in the probe request; all APs that receive the probe-request respond with a probe-response for each SSID that it supports.
- **Passive scan:** Passive scanning occurs when the client changes its 802.11 radio to the channel that is being scanned, and waits for a periodic beacon from any APs on that channel. By default, APs send beacons every 100 ms.

  Most clients use active scan because it takes 100 ms to receive a periodic beacon broadcast in a passive scan.

During a channel scan, the client is unable to transmit or receive client data traffic. Clients use the following approaches to minimize this impact to client data traffic:

- **Background scanning:** Clients scan the available channels before they roam. The scans provide information about the RF environment and available APs that can help clients to roam faster, if necessary. The affect to client traffic can be minimized by scanning only when the client is not actively transmitting data, or by periodically scanning only a single alternate channel at a time. Scanning a single channel incurs minimal data loss.

- **On-roam scanning:** On-roam scan occurs after the client determines a roam is necessary. Each vendor or device can implement its own algorithms to minimize the roam latency and the affect to data traffic. For example, some clients might scan only the nonoverlapping channels.

Typical scanning behavior

Although most client roaming algorithms are proprietary, it is possible to generalize the typical behavior. Typical wireless client roam behavior consists of the following activities:

- **On-roam scan:** Ensures that clients have the most up-to-date information at the time of the roam.
- **Active scan:** An active scan is preferred over a passive scan because of lower latency when roaming.

WLAN clients can use the following informational attributes to dynamically alter the roam algorithm:

- Client data type, for example, voice call in progress.
- Background scan information that is obtained during routine periodic background scans.

The different ways in which a WLAN client can use the attributes to alter the scan algorithm are as follows:

- **Scan a subset of channels:** For example, the client can use information from the background scan to determine channels that are being used by APs in the vicinity.
- **Terminate the scan early:** For example, if a voice call is in progress, the client can use the first acceptable AP instead of waiting to discover all APs on all channels.
- **Change scan timers:** For example, if a voice call is in progress, the client can use active scan to minimize the time that it spends waiting for probe responses.

Cisco Compatible Extensions channel scanning

While WLAN clients ultimately determine when to associate (or reassociate) to an AP, Cisco APs provide information to clients to facilitate AP selection by providing information (such as channel load in its beacons and probe responses) or by providing a list of neighboring APs.

WLC software Release 4.0 and later support the following Cisco Compatible Extensions, layer 2 client-roaming enhancements:

- **AP assisted roaming:** This feature helps clients to save scan time. Whenever a Cisco Compatible Extensions v2 client associates with an AP, it sends an information packet to the new AP listing the characteristics of its previous AP. The AP uses this information to build a list of previous APs, which it sends (through unicast) to clients immediately after association to reduce roaming time. The AP list contains the channels, basic service set identifiers (BSSIDs) of neighbor APs that support the current SSIDs of the client, and time elapsed since disassociation.
- **Enhanced neighbor list:** This feature is an Cisco Compatible Extension v4 enhancement to the neighbor list that is sent as part of the v2 AP assisted roaming feature. It is always provided automatically by the AP to the client immediately following a successful association or reassociation. Because the AP periodically checks to ensure its neighbor list is up to date, it can also send an automatic update to the corresponding clients. The enhanced neighbor list includes, for each AP, the RF parameters that are discussed in Cisco Compatible Extensions client roam triggers, on page 95. In addition, it can also include, for each AP in the list, additional information about AP timing parameters, information about the AP support for the clients subnet, and the strength and signal-to-noise ratio (SNR) of the last transmission from the client that is received by the AP.

- **Enhanced neighbor list request (E2E)** The end-to-end (E2E) specification is a Cisco and Intel joint program that defines new protocols and interfaces to improve the overall voice and roaming experience. It applies only to Intel clients in a Cisco Compatible Extensions environment. Specifically, it enables Intel clients to request a neighbor list at anytime. When this occurs, the AP forwards the request to the WLC. The WLC receives the request and replies with the current Cisco Compatible Extensions roaming sublist of neighbors for the AP to which the client is associated.

Note To check whether a particular client supports E2E, click **Wireless** > **Clients** on the WLC GUI, and then click the **Detail** link for the desired client. Also check the E2E **Version** field under **Client Properties**.

- **Directed roam request:** This feature enables the WLC to send directed roam requests to the client in situations when the WLC can better service the client on an AP that is different from the one to which the client is associated. In this case, the WLC sends the client a list of the best APs that it can join. The client can either respond to or ignore the directed roam request. Non-Cisco Compatible Extensions clients and clients running Cisco Compatible Extensions Version 3 or prior must not take any action. No configuration is required for this feature.

WLC software Release 4.0 supports Cisco Compatible Extensions Versions 1 through 4. Cisco Compatible Extensions support is enabled automatically for every WLAN on the WLC and cannot be disabled. The WLC stores the Cisco Compatible Extensions version of the client in its client database and uses it to generate and respond to Cisco Compatible Extensions frames appropriately. Clients must support Cisco Compatible Extensions Version 4 (or Cisco Compatible Extensions Version 2 for AP assisted roaming) to utilize roaming enhancements.

Many smartphones and tablets and other mobile devices are not CCX-aware and therefore do not use these CCX parameters.

Evaluating the list of potential roam targets

After the wireless client receives a list of potential APs to which it can roam, the client uses a client-specific algorithm to choose a specific AP to which it will roam. The roaming algorithm must consider the following factors in its calculations:

- Received signal strength indicator (RSSI)
- Signal-to-noise ratio (SNR)
- Number of clients on the AP
- Transmit and receive bandwidth that is being used by the AP
- RF channel load information from beacon and probe responses that is sent by the AP

Reauthenticating to a new access point

When a wireless client initially joins a WLAN, it must authenticate before it is granted access to the network. This section describes the following considerations and processes:

- Authentication types
- Reauthenticating when roaming

Authentication types

You can use the following authentication schemes for WLAN access:

- **Open authentication:** This is a null authentication; any client is permitted to access the WLAN.
- **Wired Equivalent Privacy (WEP) shared key (static WEP):** The static WEP requires both sender and receiver to have the same preprovisioned key to decode messages from each other.
- **Wi-Fi Protected Access (WPA)-Personal and WPA2-Personal:** A shared key, which is not the encryption key, is configured on both the WLAN and the WLAN client, and this key is used in the WPA four-way handshake to generate a per-session encryption key.
- **IEEE 802.1X/Extensible Authentication Protocol (EAP) authentication used in WPA-Enterprise or WPA2-Enterprise:** Depending on the deployment requirements, you can use one of the following EAP authentication protocols for secure wireless deployments:
  - Protected EAP (PEAP)
  - EAP-Transport Layer Security (EAP-TLS)
  - EAP-Flexible Authentication through Secure Tunneling (EAP-FAST)

  Regardless of the protocol that you use, all the preceding protocols currently use IEEE 802.1X, EAP, and remote authentication dial-in user service (RADIUS) as their underlying transport. Based on the successful authentication of the WLAN client, these protocols allow network access, and vitally, allow the WLAN network to be authenticated by the user.

The basic flow of an IEEE 802.1X/EAP authentication is shown in Figure 44: EAP Protocol Flow, on page 85. In that figure, the portion labeled *Authentication conversation is between client and Authentication Server* represents the authentication process between the client and the authentication server. This authentication requires the WLC to transmit multiple packets between the client and the authentication server. This portion of the authentication flow also requires CPU-intensive cryptographic processing at both the client and the authentication server. This part of the authentication is where latency can easily exceed one second and is the focus of the fast roaming algorithms that are discussed in the following section.

Reauthenticating when roaming

This section describes roaming with different authentication types:

- Roaming with open authentication or static WEP
- Roaming with IEEE 802.1X or EAP authentication
- Fast secure roaming
- Fast roaming with Proactive Key Caching

Roaming with open authentication or static WEP

When a client roams using open authentication (no keys) or using shared keys, authentication adds little roam latency. This is because no additional packets need to be exchanged between the client and the AAA server.

Roaming with IEEE 802.1X or EAP authentication

When a client roams using IEEE 802.1X with dynamic WEP, WPA-enterprise, or WPA2-enterprise, an IEEE 802.1X authentication generally must occur with an AAA/RADIUS server. Authenticating with an

AAA/RADIUS server can take more than one second. A one second interruption to latency-sensitive applications, such as voice or video over IP, when roaming is unacceptable, and therefore fast secure roaming algorithms that are developed help to reduce the roam latency.

Fast secure roaming

Fast roaming algorithms include Cisco Centralized Key Management (CCKM), Opportunistic Key Cache (OKC), and Proactive Key Caching (PKC). CCKM and PKC allow a WLAN client to roam to a new AP and reestablish a new session key, namely pairwise transient key (PTK), between the client and AP without requiring a full IEEE 802.1X or EAP reauthentication to a AAA/RADIUS server.

Both CCKM and PKC are Layer 2 roaming algorithms and therefore they do not consider any Layer 3 issues, such as IP address changes. In the Cisco Unified Wireless Network, the subnets are responsible to allocate IP addresses that originate at the WLC to the clients, and not the AP. You can achieve the following benefits from CCKM and PKC:

- Helps to group large number of WLAN clients for a given SSID into the same Layer 2 subnet.
- Maximizes the scope of the Layer 2 domain-and the fast secure roaming domain.

In addition, multiple-WLC deployments support client roaming across APs managed by WLCs in the same mobility group on the same or different subnets. This roaming is transparent to the client because the session is sustained and a tunnel between the WLCs allows the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active.

Fast secure roaming with Cisco Centralized Key Management

CCKM is a Cisco standard supported by Cisco Compatible Extensions clients to provide fast secure roaming.

CCKM requires support in the client. Cisco Compatible Extensions provides client-side specifications for support of many client functions, including fast secure roaming. The following table summarizes the supported EAP types in each version of Cisco Compatible Extensions.

Table 13: Cisco Compatible Extensions EAP Support

Cisco Compatible Extensions Version	Supported EAP Types
Cisco Compatible Extensions Version 2	CCKM with Lightweight Extensible Authentication Protocol (LEAP)
Cisco Compatible Extensions Version 3	CCKM with LEAP, EAP-FAST
Cisco Compatible Extensions Version 4	CCKM with EAP, EAP-FAST, EAP-TLS and LEAP

CCKM establishes a key hierarchy upon initial WLAN client authentication and uses that hierarchy to quickly establish a new key when the client roams. The following sections describe the initial establishment and roam phases:

- CCKM roaming - initial key hierarchy establishment
- CCKM roaming - client roam

CCKM roaming - initial key hierarchy establishment

Figure 47: CCKM Initial Key (Part 1 of 4), on page 101 through Figure 50: CCKM Initial Key (Part 4 of 4), on page 102 illustrate the initial key hierarchy establishment process. In WPA-Enterprise and WPA2-Enterprise, the outcome of a successful EAP authentication (see Figure 44: EAP Protocol Flow, on page 85) is a pairwise master key (PMK).

The following figure shows the establishment of PMK at the client and the AAA/RADIUS server, and the subsequent forwarding of the PMK to the WLC.

Figure 47: CCKM Initial Key (Part 1 of 4)

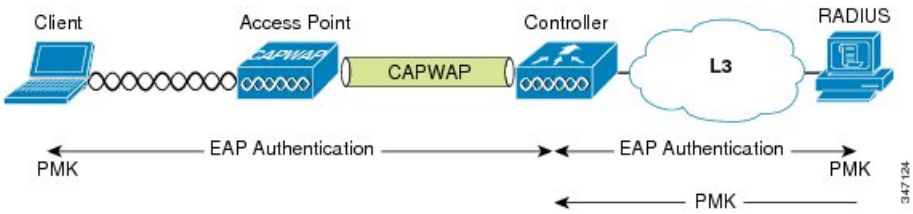

The WLC and the client both derive a network session key (NSK) from the PMK. After the NSK is established, the WPA-prescribed four-way handshake is performed between the client and the WLC. At the conclusion of the four-way handshake, a base transient key (BTK) and key request key (KRK) are established. See the following figure.

Figure 48: CCKM Initial Key (Part 2 of 4)

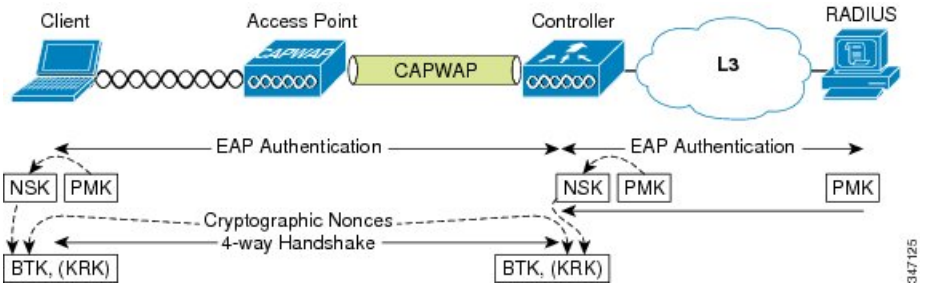

WPA and WPA2 differ only slightly from CCKM at this point. WPA/WPA2 uses the PMK directly (instead of deriving an NSK), and after the four-way handshake, establishes a pairwise transient key (PTK), thus concluding the establishment of the WPA/WPA2 unicast key.

Both the client and the WLC hash the BTK, an initial rekey number (RN)=1, and the BSSID to derive a PTK. The WLC then forwards the PTK to the AP over the Control and Provisioning of Wireless Access Points (CAPWAP) tunnel. See the following figure.

Figure 49: CCKM Initial Key (Part 3 of 4)

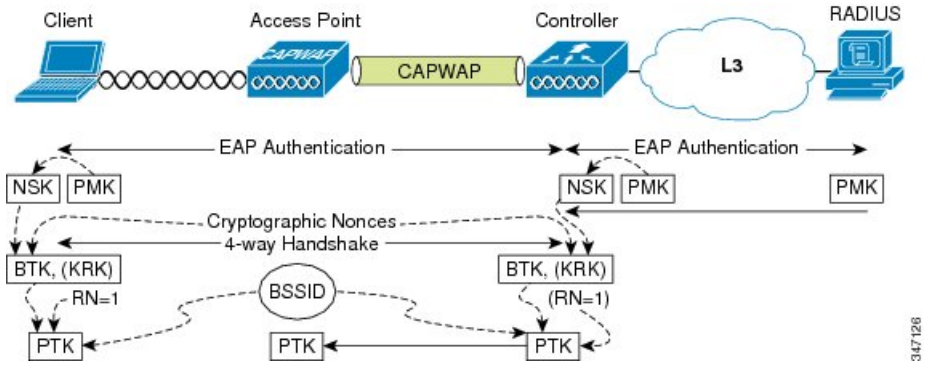

The client and AP communicate using the PTK to encrypt the data sent between them. See the following figure.

Figure 50: CCKM Initial Key (Part 4 of 4)

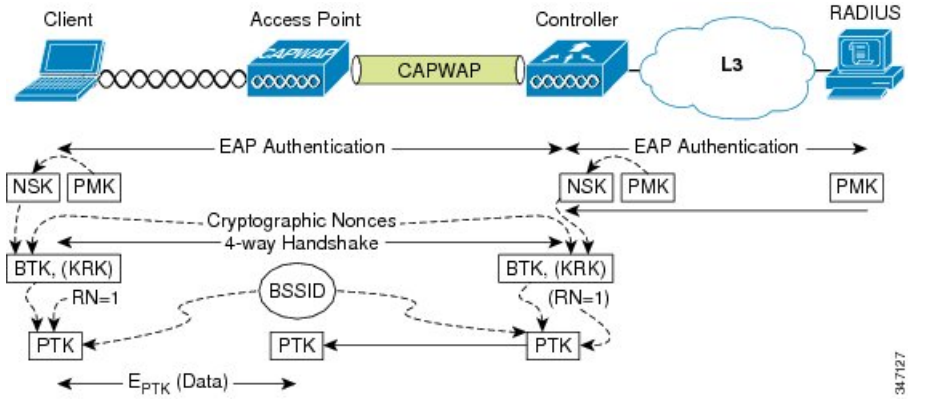

CCKM roaming - client roam

CCKM is intended to provide very fast roaming. In the absence of CCKM, a WPA/WPA2 client must perform a full EAP authentication to a remote AAA/RADIUS server, followed by a WPA/WPA2 four-way handshake whenever it roams. This process can take more than one second. With CCKM, the roaming client and WLC can use preestablished keying material to immediately establish a PTK, normally within a few tenths of a millisecond.

When the client roams to a new AP, the client sends a reassociate request with the next sequential rekey number. Protection against spoofed reassociate requests is provided by the Message Integrity Check (MIC)

that the client adds to the reassociate request (the MIC is generated using the KRK as cryptographic input). The reassociate request is forwarded by the AP to the WLC and the MIC is validated. See the following figure.

Figure 51: CCKM Roam Key (Part 1 of 2)

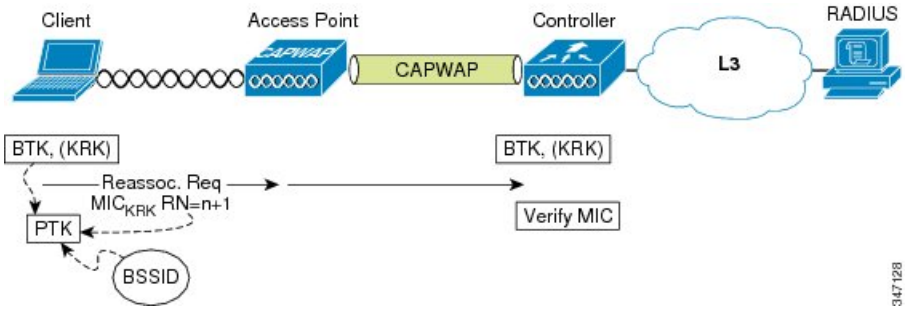

The WLC calculates the next PTK, and forwards it to the AP. The client and the AP can now communicate using the new PTK to encrypt the data sent between them. See the following figure.

Figure 52: CCKM Roam Key (Part 2 of 2)

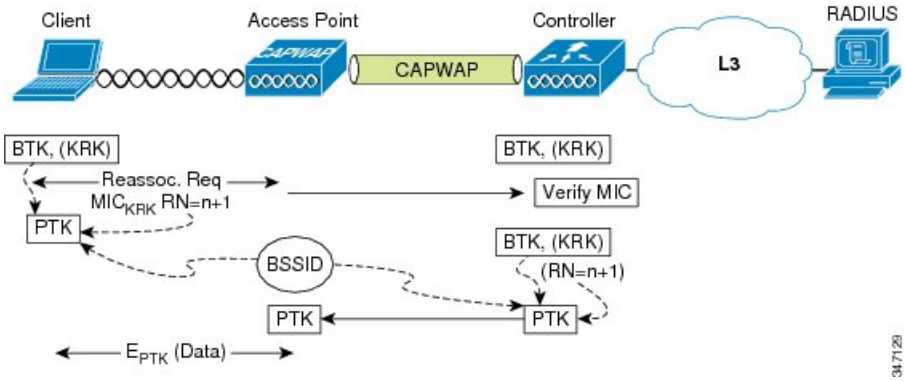

Fast roaming with proactive key caching

PKC is an IEEE 802.11i extension that allows for proactive caching (before the client roaming event) of the WPA/WPA2 PMK that is derived during a client IEEE 802.1 x/EAP authentication at the AP. See the following figure.

Figure 53: PKC Roam

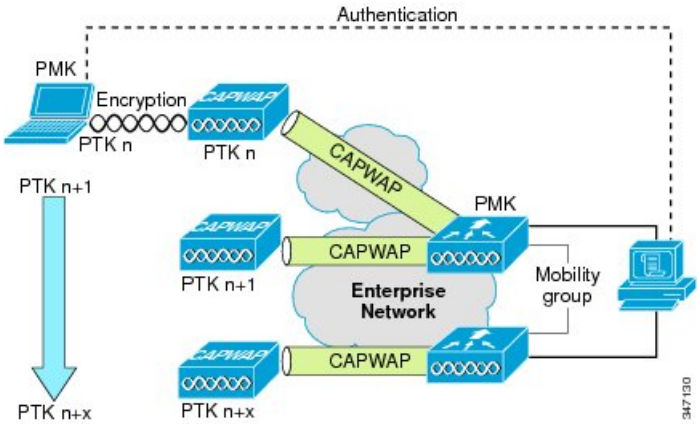

If a PMK (for a given WLAN client) is already present at an AP when presented by the associating client, full IEEE 802.1X/EAP authentication is not required. Instead, the WLAN client can simply use the WPA four-way handshake process to securely derive a new session encryption key for communication with that AP.

Note PKC is an IEEE 802.11i extension and so it is supported in WPA2, but not in WPA.

In Cisco Unified Wireless deployment, the distribution of the cached PMKs to APs is simplified. The PMK is cached in the WLCs and made available to all APs that connect to that WLC, and between all WLCs that belong to the mobility group of that WLC in advance of a client roaming event.

802.11r Fast Transition roaming

802.11r secure roaming is achieved with the exchange of fewer packets due to caching on the clients, APs, and WLC. The client preauthenticates to the *roam to AP* before the client actually roams to the *roam to AP*. So, the actual roams occurs faster because fewer packets are exchanged between the AP and the client. The packets that are exchanged happen while the client is still associated to the *roam to AP*, therefore no time is lost in the exchange of data packets, because the client is reauthenticated to the *roam to AP*.

Following are the two options for 802.11r roam configuration:

- Fast Transition (FT) roaming only over the air
- FT roaming with authentication packets on the infrastructure

The 802.11r Fast Transition authentication request and response can occur over the Wi-Fi channel as shown in the following figure.

Figure 54: 802.11r Fast Transition Roaming over Wi-Fi Channel

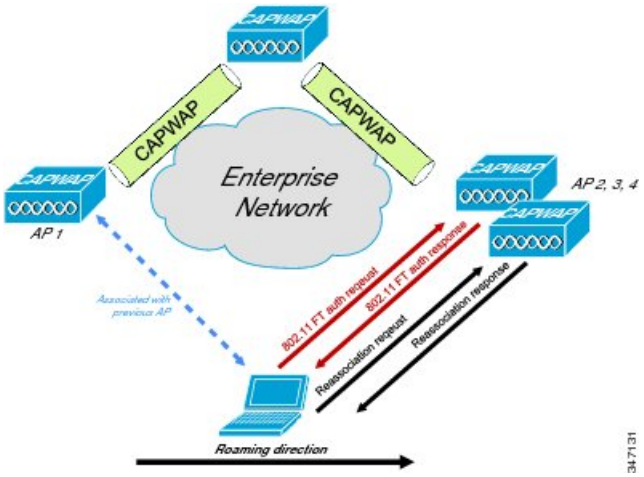

Alternatively, the 802.11r Fast Transition authentication request and response can occur over the wired network subnet. This is also known as roaming over the distributed system (DS), as shown in the following figure.

Figure 55: 802.11r Fast Transition Roaming with the Aid of the Wired Network Subnet

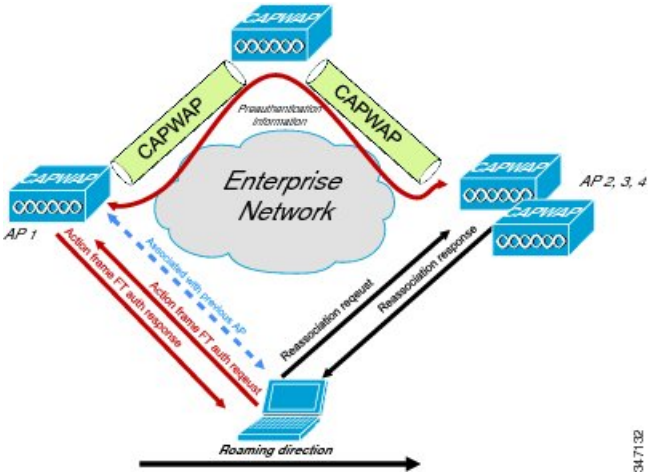

Related Topics

Cisco Compatible Extensions
Enterprise Mobility Design Guide

IP layer configuration

When a client roams from one AP to another, it must determine if it requires a new IP address, or if it can continue to use its old IP address. The client must take the following actions while roaming:

- Acquire a valid IP address through DHCP
- Enable IP duplicate address detection
- Enable Mobile IP signaling (if required)
- Virtual private network (VPN) internet key exchange (IKE) signaling (if required)

In a Cisco WLC deployment, client IP addresses do not change when they roam within the same mobility group. WLC deployments support client roaming across APs that are managed by one or more WLCs in the same mobility group on the same or different subnets. This roaming is transparent to the client because the session is sustained and a tunnel between the WLCs allow the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active.

Clients that roam without a Cisco fast secure roaming protocol (CCKM or PKC), send a DHCP request asking for their current IP address. In a Cisco WLC environment, the WLC infrastructure ensures that the client stays on the same subnet and can continue to use its old IP address. Next, the client performs duplicate address detection by pinging its own IP address to ensure that the WLAN client does not respond with the same IP address that it is using. If a client is running a mobile IP or VPN, those protocols would run after the IP address is verified unique.

Infrastructure impacts of client roaming

When a wireless client authenticates and associates with an AP, the WLC of the AP places an entry for that client in its mobility database. This entry includes the client MAC and IP addresses, security context and associations, QoS context, WLAN, and associated AP. The WLC uses this information to forward frames and manage traffic to and from the wireless client.

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

Multiple-WLC deployments support client roaming across APs managed by WLCs in the same mobility group on the same or different subnets. This roaming is transparent to the client because the session is sustained and

a tunnel between the WLCs allow the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active. The following figure illustrates the roaming in this context.

Figure 56: WLAN Infrastructure-Roam

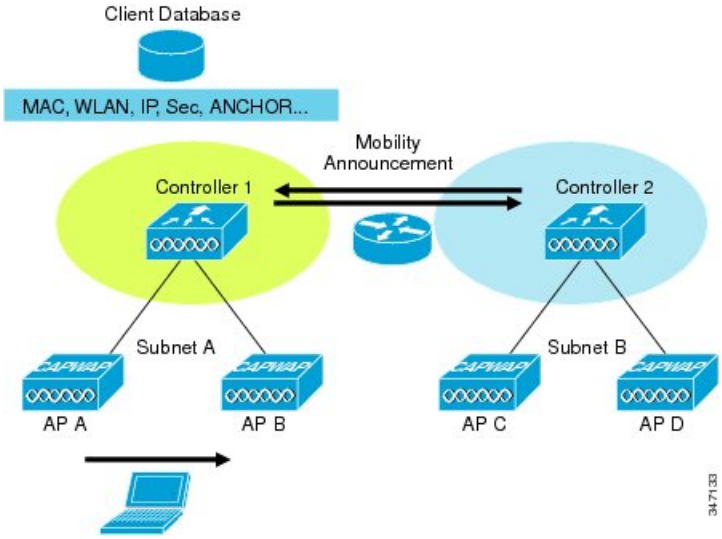

Measuring roam latency

You can segment a roam into the following components:

- Client roam decision
- Choosing a new AP to which a client roams
- Reauthenticating to the new AP
- IP layer configuration
- Infrastructure impacts of client roam

Each of the preceding components contributes to add latency to a roam. However, there is no industry consensus on how to measure roam latency. The most realistic measure of roam latency is from the last packet that is sent by the roaming client on the old AP to the first packet that is received by the roaming client on the new AP. This ensures all the preceding components are measured and ensures a two-way communication is established as illustrated in the following table:

Table 14: Summary of Roam Latency Measurement Process

Roam Action	Measurement Point	Description
Start	Last packet that is sent by roaming client on old AP	Ensures two-way communication is still established when the roam latency measurement starts. It is common for the frames to continue to be forwarded to the roaming client on the old AP after the client has started the roam.

Roam Action	Measurement Point	Description
End	First packet that is received by roaming client on new AP	Ensures two-way communication by ensuring that the clients new location has been learned by the network infrastructure and that the client is receiving packets as well as sending them.

Note When comparing roam latency for different WLAN implementations, make sure that you use the same criteria to measure roam latency in each case.

Monitoring client roaming

In addition to the Cisco Compatible Extensions Version 4 channel-scanning capabilities, Cisco Compatible Extensions Version 4 clients also send a *Roam Reason Report* to indicate why they roamed to a new AP. It also allows network administrators to build and monitor a roam history.

Use the Cisco Wireless LAN Controller command line interface commands that are listed in the following table to view information about Cisco Compatible Extensions Layer 2 client roaming:

Table 15: Cisco Compatible Extensions Layer 2 client roaming

To	Enter Command	Information Retrieved
View the current RF parameters that are configured for client roaming for 802.11a or 802.11b/g network	**show {802.11a | 802.11bg} l2roam rf-params**	Current RF parameters that are configured for client roaming for 802.11a or 802.11b/g network
View the Cisco Compatible Extensions Layer 2 client roaming statistics for a particular AP	**show {802.11a | 802.11bg} l2roam statistics ap_mac**	Use this command to have the following information: • Number of roam reason reports that are received • Number of neighbor list requests that are received • Number of neighbor list reports that are sent • Number of broadcast neighbor updates that are sent

To	Enter Command	Information Retrieved
View the roaming history for a particular client	**show client roam-history client_mac**	Use this command to have the following information: • The time when the report was received • The MAC address of the AP to which the client is currently associated • The MAC address of the AP to which the client was previously associated • The channel of the AP to which the client was previously associated • The SSID of the AP to which the client was previously associated • The time when the client disassociated from the previous AP • The reason for the client roam
Obtain debug information for the Cisco Compatible Extensions Layer 2 client roaming	**debug l2roam {detail | error | packet | all} enable**	Debug reports for the Cisco Compatible Extensions Layer 2 client roaming

802.11k management frame format

The following figure shows a decoded 802.11k neighbor list packet that is captured from a WildPackets sniffer trace. This packet was sent from the AP that an Apple iPhone 5 was associated to. The iPhone sent a neighbor request frame to the AP. The AP responded with a list of the APs that are its current neighbors. Embedded in the 802.11k neighbor list response frame is the MAC addresses of three neighbor APs and the Wi-Fi channel of each of those APs. With this information available, the 802.11k mobile client does not need to scan all the 5 GHz channels looking for candidate AP to roam to. The software of the 802.11k client can roam APs that have matching credentials and are known to be in the coverage area of the client, which saves battery life, reduces unnecessary usage of the Wi-Fi channel, and keeps the phone on the Wi-Fi channel that is doing the call processing to maintain a high-quality call with a higher mean opinion score (MOS) value. The following

figure shows all the element information that is requested by the mobile client. The 802.11k specification allows for more elements of information and more detail.

Figure 57: 802.11k Decoded Packet with Neighbor List

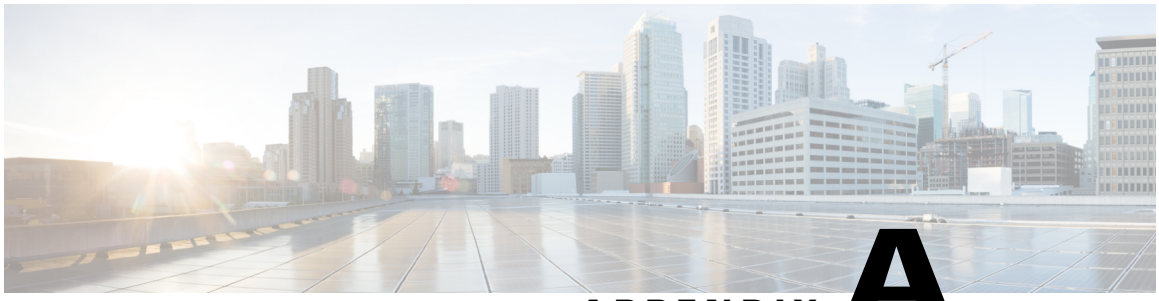

APPENDIX A

Glossary

Glossary

A

AAA	Authentication, Authorization, and Accounting
AC	Access Categories
ACM	Admission Control Mandatory
ADDTS	Add Traffic Stream
AES	Advanced Encryption Standard
ANonce	Authenticator Nonce
AP	Access Points
AVC	Application Visibility and Control
AVVID	Architecture for Voice, Video, and Integrated Data

B

BPSK	Binary Phase Shift Keying
BSS	Basic Service Set
BSSID	Basic Service Set Identifier

BT	Bluetooth
BTK	Base Transient Key
BYOD	Bring Your Own Device

C

CAC	Call Admission Control
CAPWAP	Control and Provisioning of Wireless Access Points
CCA	Clear Channel Assessment
CCKM	Cisco Centralized Key Management
CCX	Cisco Compatible Extensions
CLI	Command Line Interface
CoS	Class of Service
CSMA-CA	Carrier Sense Multiple Access-Collision Avoidance

D

dBM	Decibels per Milliwatt
DCF	Distributed Coordination Function
DFS	Dynamic Frequency Selection
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Services
DS	Distributed System
DSCP	Differentiated Services Code Point
DTIM	Delivery Traffic Indication Map

E

EAP	Extensible Authentication Protocol

EAP-FAST	EAP-Flexible Authentication via Secure Tunneling
EAP-TLS	EAP-Transport Layer Security
EAPOL	EAP over LAN
EDCF	Enhanced Distributed Coordination Function
EF	Expedited Forwarding
EMI	Electromagnetic interference
EIGRP	Enterprise Interior Gateway Routing Protocol
EOSP	End-of-Service Period

F

FMS	Fixed Mobile Substitution
FT	Fast Transition

G

GTC	Generic Token Card
GTK	Group Temporal Key
GUI	Graphical User Interface

H

HA	High Availability
HCCA	Hybrid Coordinated Channel Access
H-REAP	Hybrid Remote Edge Access Point
HSRP	Hot Standby Router Protocol
HT	High Throughput

I

IE	Information element
IKE	Internet Key Exchange
IM	Instant Messaging
IP	Internet Protocol

K

KRK	Key Request Key

L

LEAP	Lightweight EAP
LBS	Location-based Service

M

MAC	Media Access Control
MCU	Media Control Unit
MIC	Message Integrity Check
MIMO	Multiple Input - Multiple Output
MOS	Mean Opinion Score
MRC	Maximum Ratio Combining
MSCHAPv2	Microsoft Challenge-Handshake Authentication Protocol Version 2

N

NBAR	Network Based Application Recognition
NSK	Network Session Key
NTP	Network Time Services

O

OFDM	Orthogonal Frequency Division Multiplexing
OKC	Opportunistic Key Cache
OSI	Open Systems Interconnection
OSPF	Open Shortest Path First

P

PAC	Protected Access Credentials
PC	Personal Computer
PEAP	Protected EAP
PHB	Per Hop Behavior
PIM	Protocol Independent Multicast
PKC	Proactive Key Caching
PKI	Public Key Infrastructure
PLCP	Physical Layer Convergence Protocol
PMK	Pair-wise Master Key
PSTN	Public Switched Telephone Network
PTK	Pair-wise Temporal Key

Q

QAM	Quadrature Amplitude Modulation
QBSS	QoS Basic Service Set
QPSK	Quadrature Phase Shift Keying
QoE	Quality of Experience
QoS	Quality of Service

R

RADIUS	Remote Authentication Dial-in User Service
RF	Radio Frequency
RMM	Remote Management Module
RRM	Radio Resource Management
RSSI	Received Signal Strength Indication
RToWLAN	Real-Time Traffic over Wireless LAN

S

SIP	Session Initiation Protocol
SLA	Service-level Agreements
SNR	Signal-to-Noise Ratio
SNonce	Supplicant Nonce
SSID	Service Set Identifiers

T

TDM	Time Division Multiplexing
TIM	Traffic Indicator Map
TKIP	Temporal Key Integrity Protocol
TLS	Transport Layer Security
TPC	Transmitter Power Control
TSPEC	Traffic Specification
TSRS	Traffic Stream Rate Set
TXOP	Transmit Opportunity

U

U-APSD	Unscheduled Automatic Power-Save Delivery
UDP	User Datagram Protocol
UNII	Unlicensed National Information Infrastructure
UP	User Priority

V

VHT	Very High Throughput
VoWLAN	Voice over Wireless LAN
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network

W

WCS	Wireless Control System
WEP	Wired Equivalent Privacy
WLC	Wireless LAN Controller
WLAN	Wireless LAN
WMM	Wi-Fi Multimedia Mode
WPA	Wi-Fi Protected Access
WPA2	Wi-Fi Protected Access 2

Glossary

INDEX

G

H

J

K

L

M

N

O

P

Q

R

S

T

U

V

W