



CHAPTER 3

WLAN Radio Frequency Design Considerations

Introduction

This chapter describes the basic radio frequency (RF) information necessary to understand RF considerations in various wireless local area network (WLAN) environments. This chapter includes information on the following topics:

- Regulatory domains and frequencies
- Understanding the IEEE 802.11 standards
- RF spectrum implementations including 802.11b/g and 802.11a
- Planning for RF deployment
- Manually fine-tuning WLAN coverage
- Radio Resource Management (RRM), also known as Auto-RF

RF Basics

This section provides a summary of regulatory domains and their operating frequencies.

Regulatory Domains

Devices that operate in unlicensed bands do not require any formal licensing process, but when operating in these bands, the user is obligated to follow the government regulations for that region. The regulatory domains in different parts of the world monitor these bands according to different criteria, and the WLAN devices used in these domains must comply with the specifications of the relevant governing regulatory domain. Although the regulatory requirements do not affect the interoperability of IEEE 802.11b/g and 802.11a-compliant products, the regulatory agencies do set certain criteria in the standard. For example, the emission requirements for WLAN to minimize the amount of interference a radio can generate or receive from another radio in the same proximity. It is the responsibility of the vendor to get the product certified from the relevant regulatory body. [Table 3-1](#) summarizes the current regulatory domains for Wi-Fi products. The main regulatory domains are FCC, ETSI, and the MKK.

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

Table 3-1 Regulatory Domains

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

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

**Note**

See the Cisco website for compliance information and also check with your local regulatory authority to find out what is permitted within your country. The information provided in [Table 3-2](#) and [Table 3-3](#) should be used as a general guideline.

Operating Frequencies

The 802.11b/g band regulations have been relatively constant, given the length of time it has been operating. The FCC allows for 11 channels, ETSI allows for up to 13 channels, and Japan allows up to 14 channels, but requires a special license to operate in channel 14.

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

802.11b/g Operating Frequencies and Data Rates

Ratified in September 1999, the 802.11b standard operates in the 2.4 GHz spectrum and supports data rates of 1, 2, 5.5, and 11 Mbps. 802.11b enjoys broad user acceptance and vendor support. 802.11b technology has been deployed by thousands of enterprise organizations, which typically find its speed and performance acceptable for their current applications.

The 802.11g standard, which was ratified in June 2003, operates in the same spectrum as 802.11b and is backward-compatible with the 802.11b standard. 802.11g supports the additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. 802.11g delivers the same 54 Mbps maximum data rate as 802.11a, but operates in the same 2.4 GHz band as 802.11b. It also provides backward compatibility with existing 802.11b devices.

[Table 3-2](#) lists the various 802.11b/g channel frequencies and specifies whether a regulatory agency allows their use in their domain. Note that not all of these frequencies are available for use in all regulatory domains.

Table 3-2 Operating Frequency Range for 802.11b and 802.11g

Channel Identifier	Center Frequency	FCC (America)	ESTI (EMEA)	TELEC (Japan)	MOC (Israel Outdoor) ¹
1	2412	X	X	X	
2	2417	X	X	X	
3	2422	X	X	X	
4	2427	X	X	X	
5	2432	X	X	X	X
6	2437	X	X	X	X
7	2442	X	X	X	X
8	2447	X	X	X	X
9	2452	X	X	X	X
10	2457	X	X	X	X
11	2462	X	X	X	X
12	2467		X	X	X
13	2472		X	X	X
14 ²	2484			X	

¹ Israel allows channels 1 through 13 indoors.

² Japan requires a special license for channel 14.

802.11a Operating Frequencies and Data Rates

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

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

802.11a provides data rates of 6, 9, 12, 18, 24, 36, 48, with a maximum data rate of 54 Mbps, though generally at shorter ranges for a given power and gain, but it has up to 23 nonoverlapping frequency channels (depending on the geographic area) compared to the three nonoverlapping channels of 802.11b/g, 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 802.11a operates is divided into several different sections. Each of the Unlicensed National Information Infrastructure (UNII) bands presented in Table 3-3 was originally intended for different uses, but all can currently be used by indoor 802.11a with appropriate power restrictions. Initially, the FCC defined only the UNII-1, UNII-2, and UNII-3 bands, each of which had four channels. The channels were spaced 20 MHz apart with an RF spectrum bandwidth of 20 MHz, thereby providing nonoverlapping channels.

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

The channels in UNII-1 (5.150 to 5.250 GHz) are 34, 36, 38, 40, 42, 44, 46, and 48. The channels in UNII-2 (5.250–5.350 GHz) are 52, 56, 60, 64 and require Dynamic Frequency Selection (DFS) and Transmitter Power Control (TPC). The channels in the new frequency range (5.470–5.725 GHz) are 100, 104, 108, 112, 116, 120, 124, 128, 132, 136, 140 and require DFS and TPC. The channels in UNII-3 are 149, 153, 157, 161, 165 and require DFS and TPC. Not all channels in a given range can be used in all of the regulatory domains. Figure 3-1 shows the various channels in the UNII-1, 2, and 3 bands, along with the additional 11 new channels.

For more information on FCC regulation updates, see the following URL:

http://www.cisco.com/en/US/products/hw/wireless/ps469/products_white_paper0900aecd801c4a88.shtml

Table 3-3 Operating Frequency Range for 802.11a

Channel Identifier	36	40	44	48	52	56	60	64	149	153	157	161
Center Frequency	5180	5200	5220	5240	5260	5280	5300	5320	5745	5765	5785	5805
Band	UNII-1				UNII-2				UNII-3			

Table 3-3 shows the standard 802.11a frequencies. Table 3-4 shows the specific frequency bands and channel numbers for a few specific regulatory domains.

Table 3-4 Additional Frequency Bands and Channel Numbers for Other Regulatory Domains

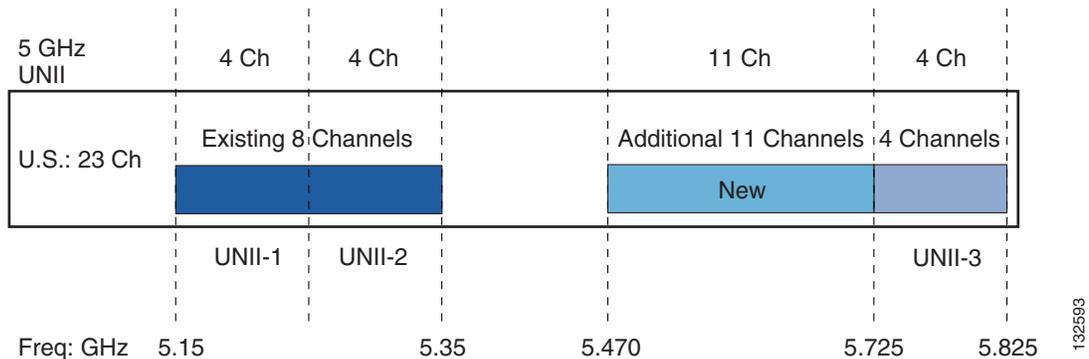
Regulatory Domain	Frequency Band	Channel Number	Center Frequency
Japan ¹	U-NII lower bands	36	5.180
		40	5.200
		44	5.220
		48	5.240
Singapore	U-NII lower band	36	5.180
		40	5.200
		44	5.220
		48	5.240
Taiwan		52	5260
		56	5280
		60	5300
		64	5320
EMEA 1 Australia New Zealand	Same as USA	Same as USA	Same as USA
EMEA 2 ²	U-NII lower band	36	5.180
		40	5.200
		44	5.220

¹ Japan is changing from channels 34,38,42,46 to 36,40,44,48 and adding channels 52, 56, 60, and 64 but required Dynamic Frequency Selection on channels 52 - 64. Cisco Equipment with the old channels are designated with a -J, Cisco equipment with the new channels are designated with a -P.

² Some EMEA countries, such as Denmark and Germany, are limited to 20 mW.

In February of 2004, the FCC released a revision to the regulations covering the 5 GHz 802.11a channel usage. This revision added 11 additional channels, bringing the available channels capacity to 23 channels (see Figure 3-1). The new additional 11 channels will be for Indoor/Outdoor use. To use the 11 new channels, however, radios must comply with two features that are part of the 802.11h specification: transmit power control (TPC) and dynamic frequency selection (DFS). DFS is required to avoid radar that operates in this frequency range, but it can also be used for other purposes, such as dynamic frequency planning. 802.11h has been supported since Cisco Unified Wireless Network Software Release 3.1.

Figure 3-1 802.11 Channel Capacity



Understanding the IEEE 802.11 Standards

IEEE 802.11 is the working group within the Institute for Electrical and Electronics Engineers (IEEE) responsible for wireless LAN standards at the physical and link layer (Layer 1 and Layer 2) of the OSI model, as compared to the Internet Engineering Task Force (IETF, which works on network layer (Layer 3) protocols. Within the 802.11 working group are a number of task groups that are responsible for elements of the 802.11 WLAN standard. Table 3-5 summarizes some of the task group initiatives.

For more information on these working groups, see the following URL:

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

Table 3-5 IEEE 802.11 Task Group Activities

Task Group	Project
MAC	To develop one common MAC for WLANs in conjunction with a physical layer entity (PHY) task group
PHY	To develop three WLAN PHYs—Infrared, 2.4 GHz FHSS, 2.4 GHz DSSS
a	To develop PHY for 5 GHz UNII band
b	To develop higher rate PHY in 2.4 GHz band
c	To cover bridge operation with 802.11 MACs (spanning tree)
d	To define physical layer requirements for 802.11 operation in other regulatory domains (countries)
e	To enhance 802.11 MAC for QoS
f	To develop recommended practices for Inter Access Point Protocol (IAPP) for multi-vendor use
g	To develop higher speed PHY extension to 802.11b (54 Mbps)
h	To enhance 802.11 MAC and 802.11a PHY-Dynamic Frequency selection (DFS), Transmit Power control (TPC)
i	To enhance 802.11 MAC security and authentication mechanisms
j	To enhance the 802.11 standard and amendments to add channel selection for 4.9 GHz and 5 GHz in Japan

Table 3-5 IEEE 802.11 Task Group Activities (continued)

k	To define RRM enhancements to provide interfaces to higher layers for radio and network measurements
k	To define Radio Resource Measurement enhancements to provide interfaces to higher layers for radio and network measurements
m	To perform editorial maintenance, corrections, improvements, clarifications, and interpretations relevant to documentation for 802.11 family specifications
n	Focus on high throughput extensions (>100MB/s at MAC SAP) in 2.4GHz and/or 5GHz bands
o	To provide Fast Handoffs in Voice over WLAN (goal is around 50ms)
p	Focus on vehicular communications protocol aimed at vehicles, such as toll collection, vehicle safety services, and commerce transactions via cars
r	To develop a standard specifying fast BSS transitions and fast roaming
s	To define a MAC and PHY for meshed networks that improves coverage with no single point of failure
t	To provide a set of performance metrics, measurement methodologies, and test conditions to enable manufacturers, test labs, service providers, and users to measure the performance of 802.11 WLAN devices and networks at the component and application level
u	To provide functionality and interface between an IEEE 802.11 access network (Hotspot) and any external network
v	To provide extensions to the 802.11 MAC/PHY to provide network management for stations (STAs)
w	To provide mechanisms that enable data integrity, data origin authenticity, replay protection, and data confidentiality for selected IEEE 802.11 management frames including but not limited to: action management frames, deauthentication and disassociation frames

RF Spectrum Implementations

In the United States, three bands are defined as unlicensed: industrial, scientific, and medical (ISM) bands. The ISM bands are as follows:

- 900 MHz (902-to-928 MHz)
- 2.4 GHz (2.4-to-2.4835 GHz) (IEEE 802.11b/g operates in this frequency range)
- 5 GHz (5.15-to-5.35 and 5.725-to-5.825 GHz) (IEEE 802.11a operates in this frequency range)

Each range has different characteristics. The lower frequencies exhibit better range, but with limited bandwidth and thus lower data rates. The higher frequencies exhibit less range and are subject to greater attenuation from solid objects.

The following sections cover some of the specific RF characteristics that the various 802.11 radios use for improving communications in the 2.4 and 5 GHz frequency ranges.

Direct Sequence Spread Spectrum

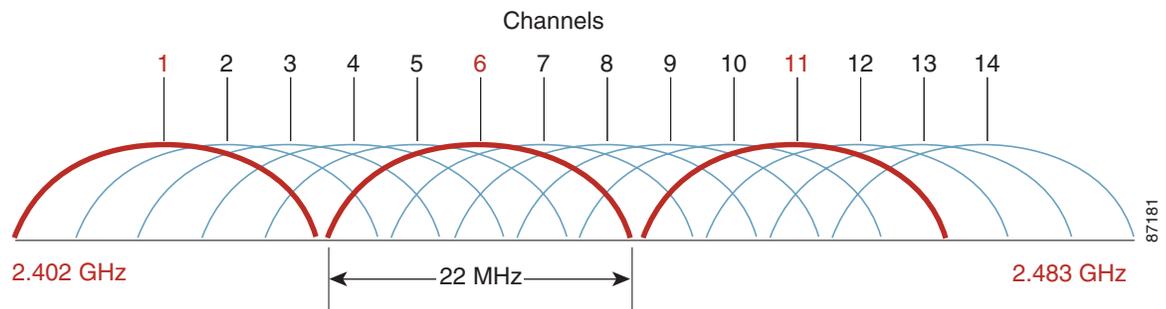
Direct sequence spread spectrum (DSSS) encodes redundant information into the RF signal. This provides the 802.11 radio with a greater chance of understanding the reception of a packet, given background noise or interference on the channel. Every data bit is expanded into a string of bits, or chips,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

IEEE 802.11b Direct Sequence Channels

14 channels are defined in the IEEE 802.11b direct sequence (DS) channel set. Each DS channel transmitted is 22 MHz wide, but the channel separation is only 5 MHz. This leads to channel overlap such that signals from neighboring channels can interfere with each other. In a 14-channel DS system (11 usable channels in the US), only three nonoverlapping (and thus, non-interfering) channels 25 MHz apart are possible (channels 1, 6, and 11).

This channel spacing governs the use and allocation of channels in a multi-AP environment, such as an office or campus. APs are usually deployed in a cellular fashion within an enterprise, where adjacent APs are allocated nonoverlapping channels. Alternatively, APs can be co-located using channels 1, 6, and 11 to deliver 33 Mbps bandwidth to a single area (but only 11 Mbps to a single client). The channel allocation scheme is illustrated in [Figure 3-2](#).

Figure 3-2 IEEE 802.11 DSS Channel Allocations



IEEE 802.11g

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

Table 3-6 lists 802.11g modulation and transmission types for the various data rates.

Table 3-6 802.11g Modulation and Transmission Types

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

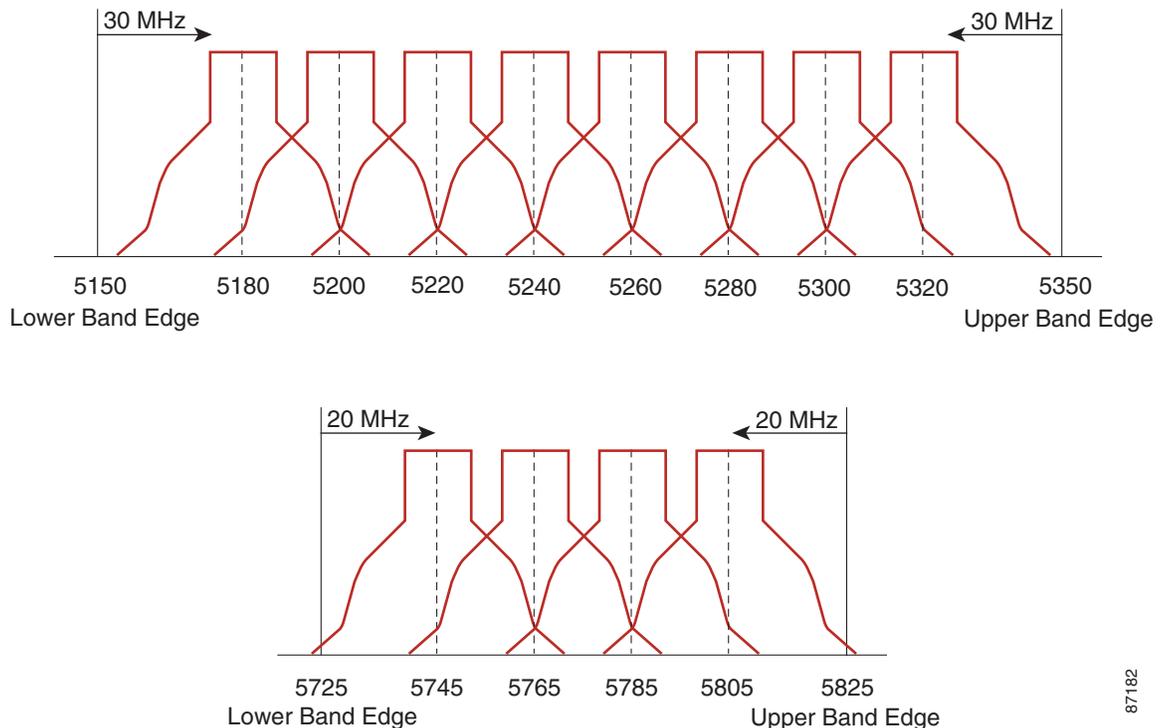
IEEE 802.11a OFDM Physical Layer

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

IEEE 802.11a Channels

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

Figure 3-3 Channel Set



For the US-based 802.11a standard, the 5 GHz unlicensed band covers 300 MHz of spectrum and supports 12 channels. As a result, the 5 GHz band is actually a conglomerate of three bands in the USA: 5.150-to-5.250 GHz (UNII 1), 5.250-to-5.350 GHz (UNII 2), and 5.725-to-5.875 GHz (UNII 3).

RF Power Terminology

RF-specific terms, such as dB, dBi, and dBm are used to describe the amount of change in power transmitted measured at any discrete point, as perceived by the radio. The following sections cover their differences and provide a rule of thumb for their use, in addition to providing an explanation of effective isotropic radiated power (EIRP).

dB

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

The math is as follows:

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

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

Remember that it is logarithmic (meaning that it increases or decreases exponentially and not linearly), and it is a ratio of some value to a reference. Also, remember that it is multiplied by 10.

Given that it is logarithmic, there are some general rules of thumb. An increase or decrease of 3 dB means that the signal doubled (double the power) or reduced in strength by 1/2, respectively. An increase or decrease of 10dB means that the signal went up by 10 times or down to 1/10th the original value.

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

dBi

The term *dBi* is used to describe the power gain rating of antennas. The real antennas are compared to an isotropic antenna (a theoretical or imaginary antenna) that sends the same power density in all directions, thus the use of dBi.

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

dBm

The term *dBm* uses the same calculation as described in the dB section, but has a reference value of 1 milliwatt.

So, taking into consideration the example previously given in the dB section, if the power jumped from 1 mW to 100mW at the radio, the power level would jump from 0 dBm to 20 dBm.

Besides describing transmitter power, dBm can also describe receiver sensitivity. Receiver sensitivity is in minus dBm (-dBm), because the signal reduces in value from its point of transmission. The sensitivity indicates the lowest power the receiver can receive before it considers the signal unintelligible. The approximate receiver sensitivity of Cisco radios is -84dBm for a 1200 series “a” radio and -90 dBm for the 1200 “g” radio.



Note

The “g” radio can recover a signal at half the strength of the “a” radio. It is generally true that for a given technology, lower frequency radios can achieve a better sensitivity than higher frequency radios.

Effective Isotropic Radiated Power

Although transmitted power based on the radio setting is rated in either dBm or Watts, the maximum energy density coming from an antenna from a complete system is measured as effective isotropic radiated power (EIRP), which is a summation of the dB values of the various components. EIRP is the value that regulatory agencies, such as the FCC or ETSI, use to determine and measure power limits, expressed in terms of maximum energy density within the first Fresnel of the radiating antenna. EIRP is calculated by adding the transmitter power (in dBm) to antenna gain (in dBi) and subtracting any cable losses (in dB). For example, if you have a Cisco Aironet 350 bridge connected to a solid dish antenna by a 50 foot length of coaxial cable, plugging in the numbers gives the following:

- Bridge—20 dBm
- 50 Foot Cable—3.3 dBm (negative because of cable loss)
- Dish Antenna—21 dBi
- EIRP—37.7dBm

Planning for RF Deployment

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

Different Deployment Types of Overlapping WLAN Coverage

How much overlapping WLAN coverage you set in your wireless network depends on the usage, though with limited exceptions, all designs should be deployed to enable low latency coverage. Wireless networks can be deployed for location management, voice, or data-only networks, or a combination of all three. The difference is in the pattern in which the APs are laid out, and the amount of RF overlap in the coverage area. When planning a WLAN deployment consideration should be given to future uses of the WLAN deployment.

Converting a WLAN deployment to support additional services beyond a data-only deployment is not simply a matter of adding APs; it can require an additional site survey and the possible relocation of existing APs.

Data-Only Deployment

Data-only deployments do not require a large amount of overlap. This is because 802.11 clients respond to a lower signal from a nearby AP, should one fail, by stepping down their rate and taking a longer time to transmit. The required overlap is determined by the WLAN data rate requirement described in [WLAN Data Rate Requirements, page 3-15](#). Minimal overlap is required for data-only networks, which allow all data rates. For data-only networks, the rule of thumb for separation of APs is typically 120–130 feet, but, when making your estimation for AP separation, remember to factor in objects that affect RF coverage, such as wall densities, machinery, elevators, or even wide-open space with steel cages, because your results can vary depending on the RF environment. RRM has been developed for this type of deployment and it is very useful for controlling the RF coverage.

Voice/Deployment

[Figure 3-4](#) shows the voice network pattern and overlap. The APs are grouped closer together and have more overlap than a data-only installation, because voice clients should roam to a better AP before dropping packets. You generally also want to run smaller cells than in the past, and ensure the overlapping cell edges at or above -67 dBm. This accomplishes a number of things including greater homogeneity across a single cell and reducing processor load in the handheld, which increases link stability and reduces latency. Although only one AP might be required for a defined area, Cisco recommends that you have two APs on nonoverlapping channels with a received signal strength indication (RSSI) above 35 at all times in your installation, for redundancy and load balancing purposes. For the 7920 voice deployment, Cisco recommends that you have a Received Signal Strength Indication (RSSI) above 35 at all time in your installation. This is to ensure that the VoIP phone has good reception as well as allowing some over-subscription and enhances roaming choices for the phone.

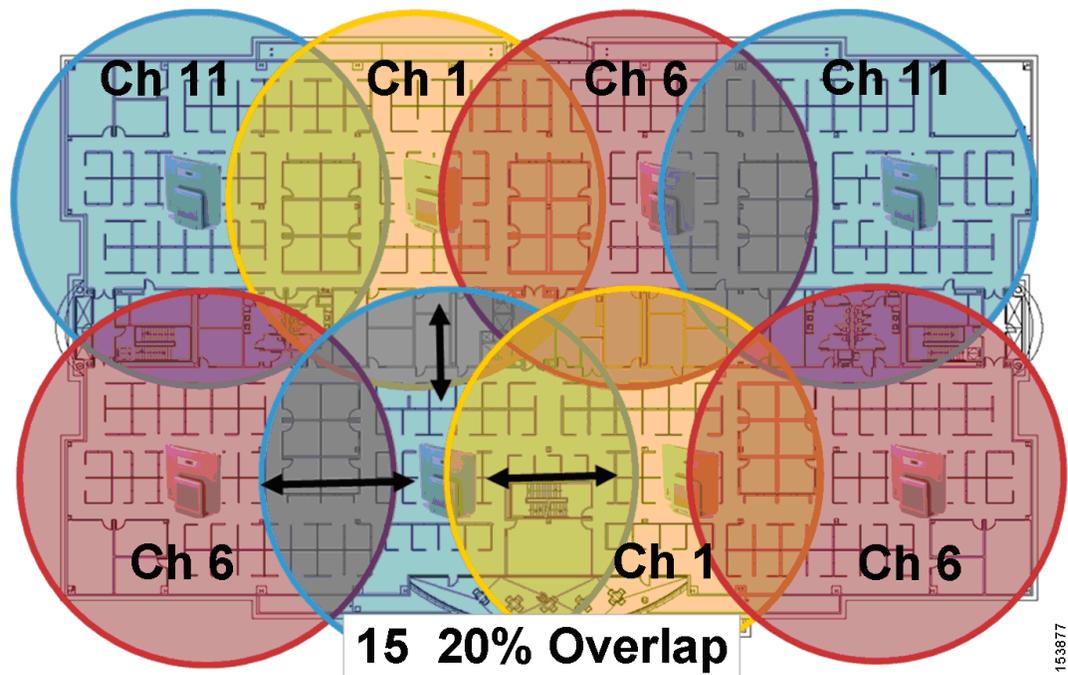
Remember that designing for low noise background is as important as relatively high energy density within the cell. This means that a good baseline power setting for the AP is in the 35–50 mW range. This generally requires approximately 15 percent more APs than if you deployed a coverage model at 100 mW. The smaller the cell, generally the more homogenous the energy density across the cell.

100 mW designs also generate significantly greater noise load into a given area; 35–50 mW designs are commonly 10 dBm or lower in background noise, even in complex coverage areas.

Pre-site surveys are useful for identifying and characterizing certain challenging areas and potential sources for interference, such as existing WLANs, rogues, and non-802.11 interference from sources such as microwave ovens and cordless telephones. Following a design that should be reviewed and approved by all stakeholders, post-site surveys should be considered as an excellent audit mechanism to ensure that the coverage model complies with the intended functional requirements as set forth by the stakeholders.

When making your estimation for separation, remember to factor in objects that affect RF coverage such as wall densities, machinery, elevators, or even wide open spaces with steel cages, because your results may vary depending on the RF environment. Be sure to include transient dynamics such as forklifts, large groups of people, or large objects moved through the area by crane or similar load bearing devices. Characterizing for other wireless deployments such as robotic systems, cranes, and so on, are also a good idea. A WLC is often a very effective method for preliminary site evaluation, by allowing a fast deployment of a WLAN infrastructure that can then be used to make RF measurements of the area; a hand-walked site survey is also effective insurance for complex areas such as those commonly found in healthcare, retail, and manufacturing.

Figure 3-4 Single Floor Site Survey for Voice



153877

For more information on a wireless voice deployment, see [Chapter 11, “VoWLAN Design Recommendations,”](#) as well as the 7920 deployment guide at the following URL:
http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7920/5_0/english/design/guide/7920ddg.html.

Location-Based Services Deployments

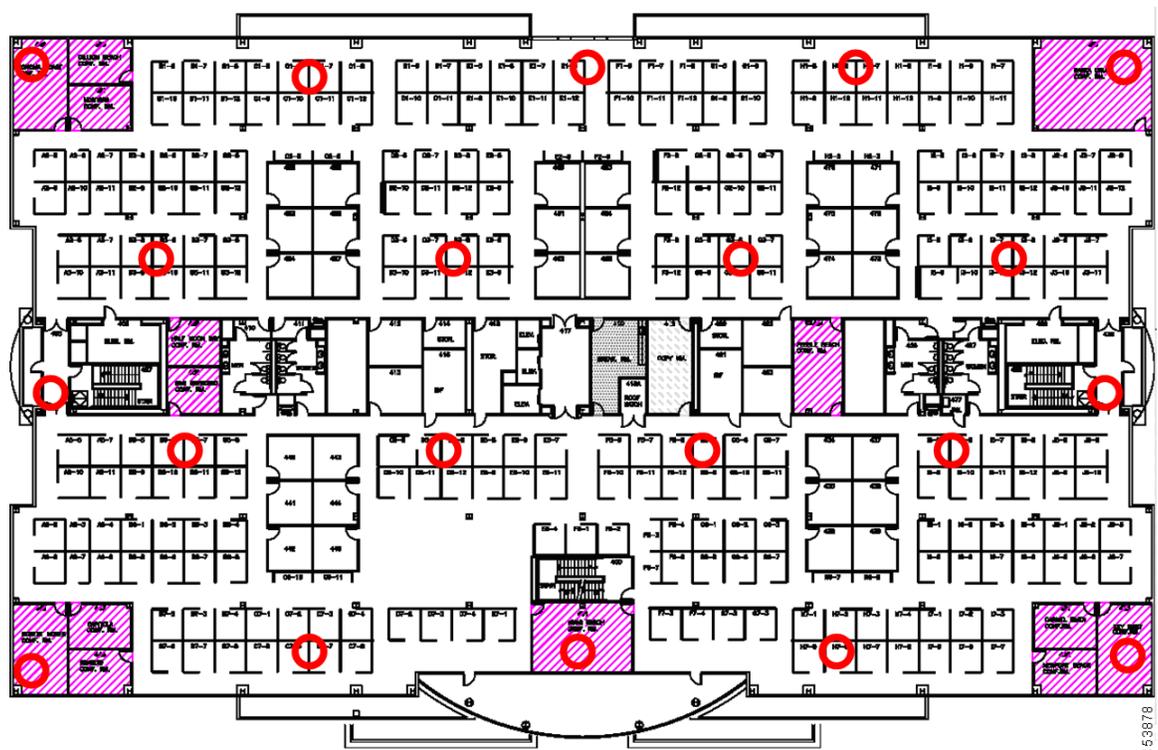
The third type of deployment is the location-based services (LBS) deployments, which may be the most complex of current applications because it relies not only on excellent cell coverage, but optimal location of APs. Location management deployments can simultaneously track thousands of devices by using the WLAN infrastructure. Examples include Wi-Fi tag type deployments or asset tracking deployments to locate equipment or devices via the wireless network and/or simply to indicate where wireless clients are throughout the wireless network in relation to a drawing or diagram. This can be used to make the wireless infrastructure more secure by providing the location of a rogue client or APs, and greatly improve client troubleshooting capabilities.

For a location management deployment, the APs are laid out in a staggered pattern. [Figure 3-5](#) shows a typical pattern. The staggered pattern allows for more accurate estimation of the location of a device.

For a discussion of Location-Based Services, see [Chapter 15, “Cisco Unified Wireless Location-Based Services,”](#) and the white paper entitled *WiFi Location Based Services 4.1 Design Guide*, which can be found at the following URL:

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

Figure 3-5 Example of a Single Floor Location Management Deployment



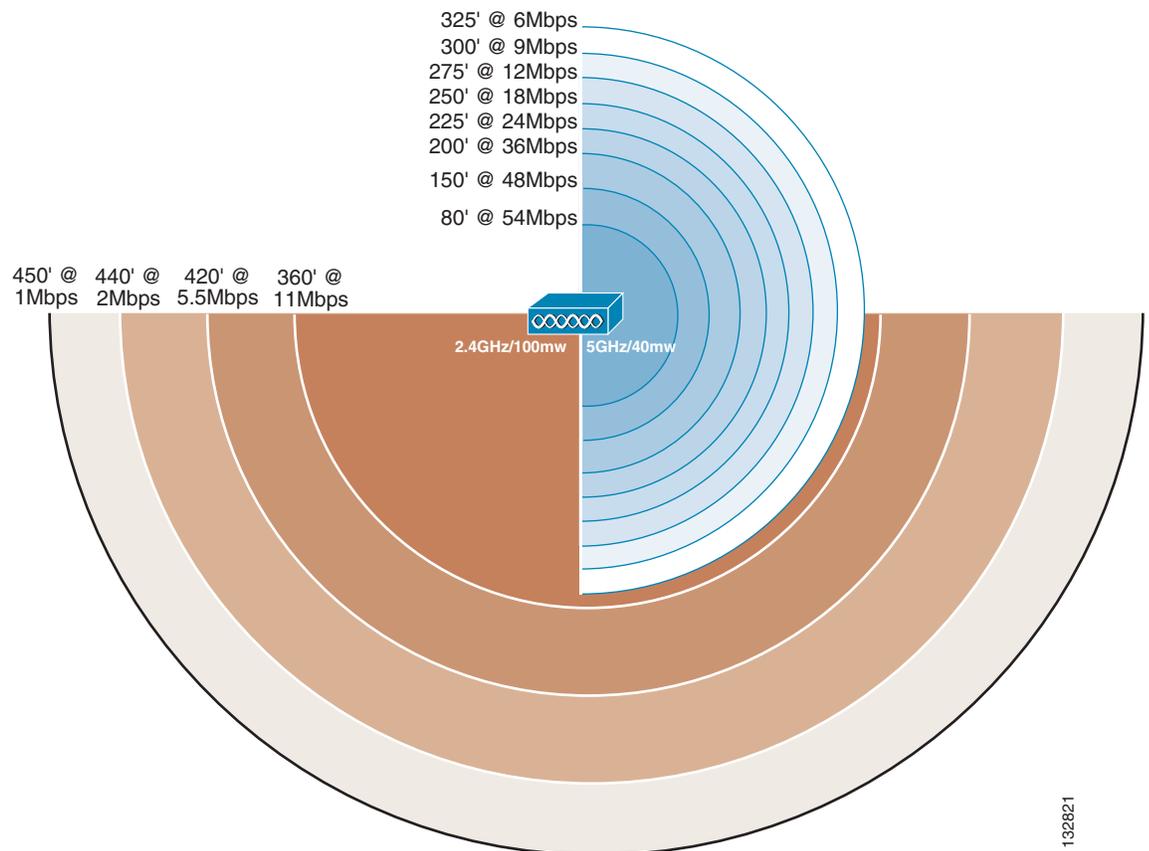
WLAN Data Rate Requirements

Data rates affect AP coverage areas. Lower data rates (such as 1 Mbps) can extend the coverage area farther from the AP than higher data rates (such as 54 Mbps) as illustrated in [Figure 3-6](#) (which is not drawn to scale). Therefore, the data rate (and power level) affects coverage and consequently the number of APs required for the installation, as illustrated in [Figure 3-7](#) for different data rates. As part of the planning process, consider the required data rates, the required range, and the required reliability.

Data Rate Compared to Coverage Area

Different data rates are achieved by the AP sending a redundant signal on the wireless link, allowing data to be more easily recovered from noise. The number of symbols, or chips, sent out for a packet at the 1 Mbps data rate is greater than the number of symbols used for the same packet at 11 Mbps. This means that sending data at the lower bit rates takes more time than sending the equivalent data at a higher bit rate. And when there is more than one client associated to the radio, the lower rate client affects the higher rate clients' maximum data rate by taking longer to transmit a packet of the same length.

Figure 3-6 Data Rate Compared with Coverage

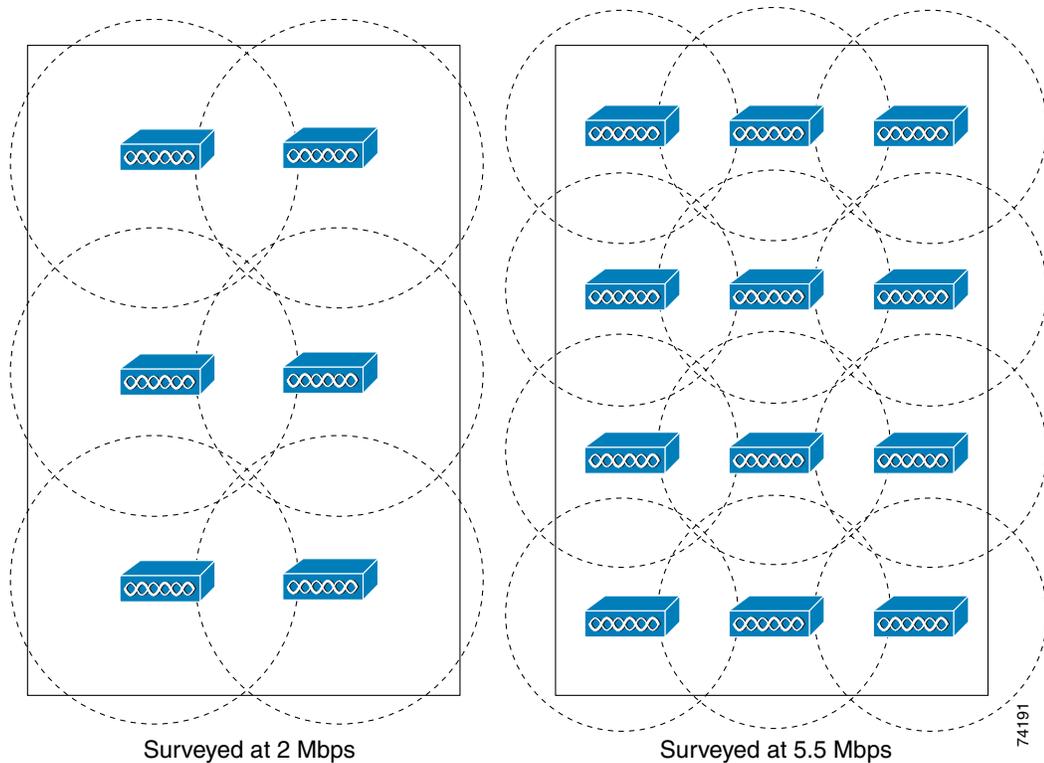


The actual diameter of the coverage depends on factors such as environment, power level, and antenna gain. For example, indoors using the standard antennas on the NIC card and APs, the diameter of the 1 Mbps circle is approximately 700 feet (210 m), and the diameter of the 11 Mbps circle is about 200 feet (60 m). Increasing the gain of the antenna can increase the distance and change the shape of the radiation pattern to something more directional.

AP Density for Different Data Rates

The minimum required reliable data rate has a direct impact upon the number of APs needed in the design, along with power setting, antenna gain, and location. Figure 3-7 shows coverage comparison and AP density for different data rates. Although six APs with a minimum data rate of 2 Mbps might adequately service an area, it might take twice as many APs to support a minimum data rate of 5 Mbps, and more again to support a minimum data rate of 11 Mbps for the same coverage area.

Figure 3-7 Coverage Comparison and AP Density for Different Data Rates



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

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

- Broadcast and multicast (if enabled) are sent at the lowest associated data rate (to ensure that all clients can receive the packets). This reduces the throughput of the WLAN because traffic must wait until frames are processed at the slower rate.

- Clients that are farther away, and therefore accessing the network at a lower data rate, decrease the overall throughput by causing delays while the lower bit rates are being serviced. It might be better to force the clients to roam to a closer AP so as not to impact the performance of the rest of the network.
- If a 54 Mbps service is specified and provisioned with APs to support *all* data rates, clients at lower rates can associate with the APs that can create a coverage area greater than planned, thereby increasing the security exposure (by allowing association from outside the building) and potentially interfering with other WLANs.

Client Density and Throughput Requirements

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

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

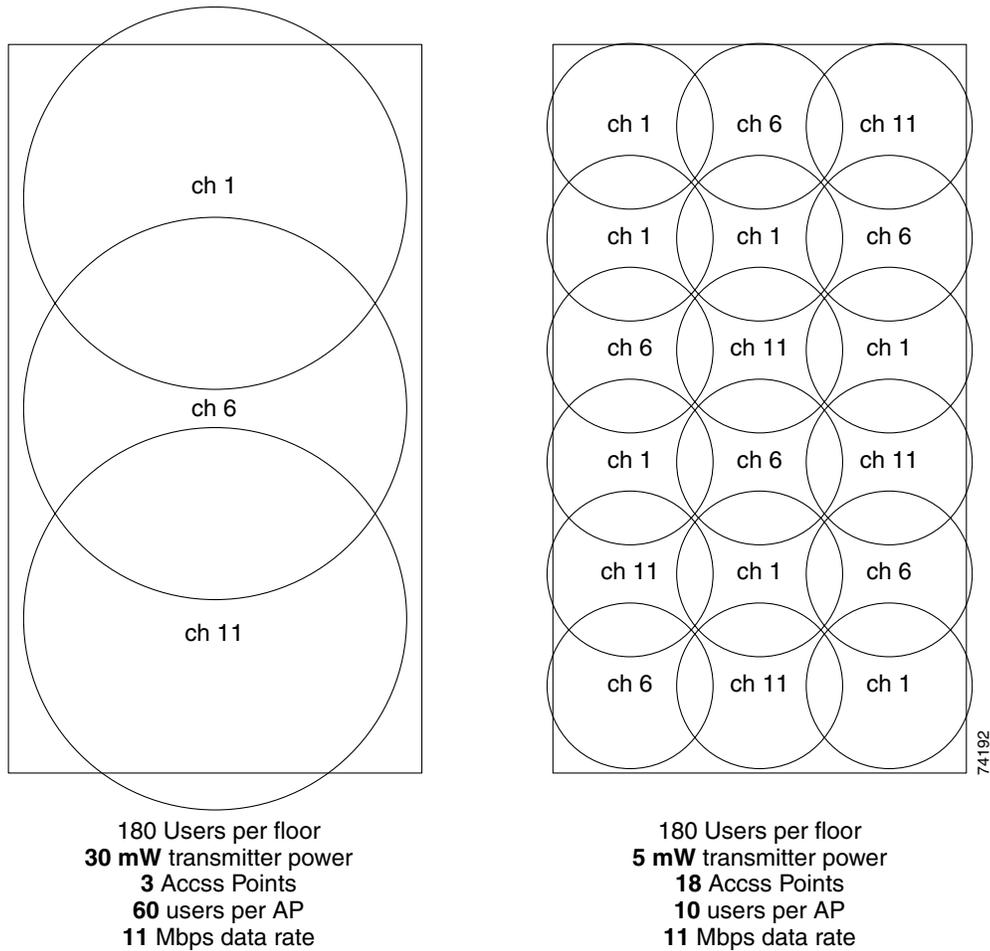
With this in mind, you must have the approximate estimate of the maximum number of active associations (active clients). This can be adjusted more or less according to the particular application.

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

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

You can increase the potential per-user throughput by decreasing the number of users contending for the aggregate throughput provided by a single AP. This can be done by decreasing the size of the coverage area, or adding a second AP on a non-overlapping channel in the same coverage area. To reduce the coverage area, the AP power or antenna gain can be reduced, resulting in fewer clients in that coverage area. This means you need more APs for the same overall area, increasing the cost of deployment. An example of this is shown in [Figure 3-8](#).

Figure 3-8 Changing the Output Power to Increase Client Performance



Note

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

WLAN Coverage Requirements

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

Power Level and Antenna Choice

Power level and antenna choice go hand-in-hand to determine AP placement. Together, these two variables determine where and how powerful the RF is in any given place in the environment. Along with choosing the correct antenna to produce the required coverage area, Cisco recommends the use of RRM to control the power level and provide the optimal channel/power plan through the use of special reuse (see [Radio Resource Management \(Auto-RF\)](#), page 3-28 for more information).

To understand antenna principles a little better, a brief introduction to the three fundamental properties follows:

- An antenna gives the wireless system three fundamental properties: gain, direction, and polarization. Gain and direction mandate range, speed, and reliability; polarization affects reliability and isolation of noise.

Gain is a measure of increase in power. Direction is the shape of the transmission pattern. A good analogy for an antenna is the reflector in a flashlight. The reflector concentrates and intensifies the light beam in a particular direction similar to what a parabolic dish antenna does to an RF source in a radio system. RF has an electric field component and a magnetic field component that are typically confined to a plane perpendicular to the propagation direction. Polarization indicates the direction of the electric field; if the electric field is orientated vertically, the wave is said to be vertically polarized. Polarization can be simply related to how the antenna is positioned. An antenna that is parallel to the Earth's surface produces horizontally polarized radio waves.

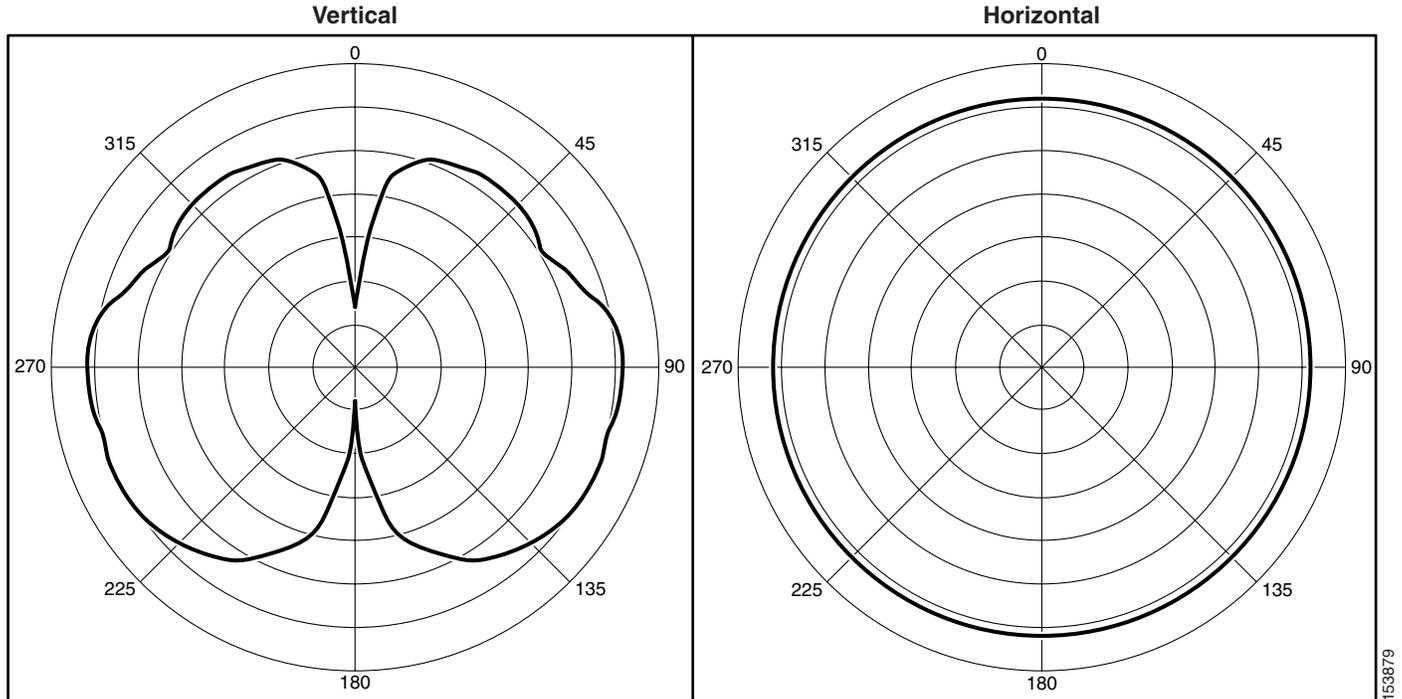
Omni and Directional Antennas

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

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

For example, with an omni-directional antenna, the antenna radiates in a donut pattern. [Figure 3-9](#) shows an omni-directional antenna RF radiation pattern in the vertical and horizontal direction. This is an actual measurement, so it does not follow the donut lines perfectly, but does show from where the donut pattern comes. As described above, other RF-affecting variables (people in the room, amount of devices stored in the facility, leaves on trees for outdoor deployment, interference from different RF sources, and so on) may affect the real RF coverage pattern.

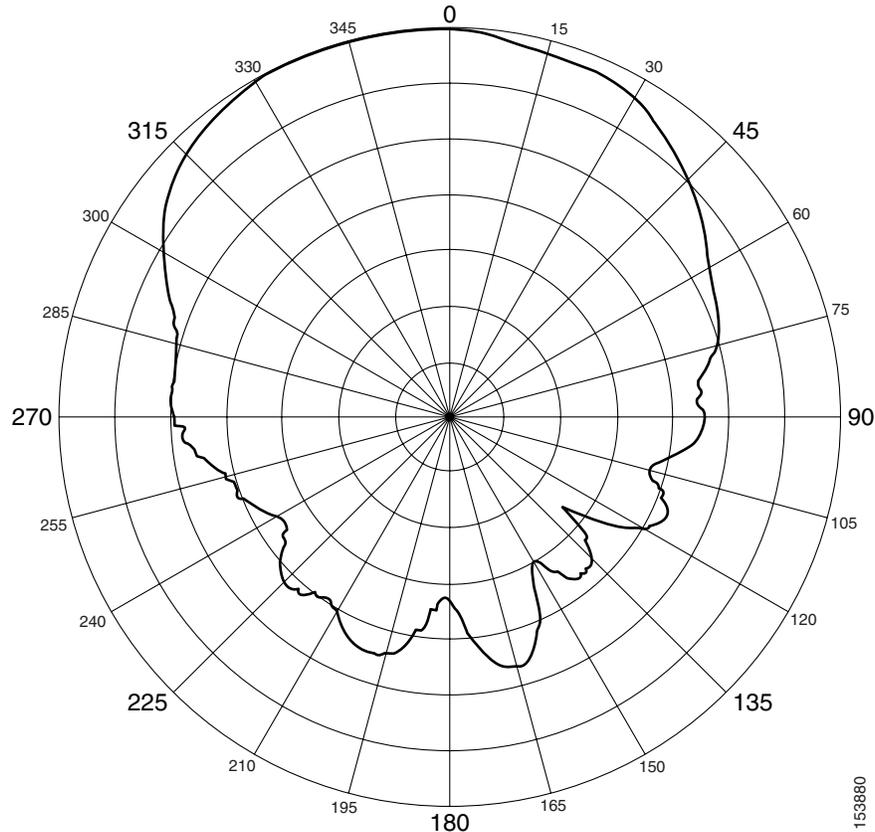
Figure 3-9 Omni-Directional RF Pattern



Looking at the pattern in [Figure 3-9](#), this may be the incorrect antenna to use on a wall, especially if it is mounted along an exterior wall where the pattern can radiate outside of the building. This can open up the wireless network to hackers outside the building and compromise the wireless network.

Patch Antennas

Another type of directional antenna is a patch antenna. Patch antennas not only radiate away from the wall or place where they are mounted, but also have rear and side lobes that produce a weakened but still quite effective RF region. [Figure 3-10](#) shows the real horizontal pattern of a diversity patch wall mount antenna. Although most of the coverage area is in front of the patch antenna, notice the back and side RF pattern from the center area. Again, antenna selection is important because it defines the radiation pattern and where wireless connectivity is possible.

Figure 3-10 Patch Wall Mount Antenna Horizontal Plane

For more information on antenna selection, see the *Cisco Antenna Selection Guide* at the following URL: http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/product_data_sheet09186a008008883b.html.

Security Policy Requirements

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

RF Environment

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

- 2.4 GHz cordless phones
- Walls fabricated from wire mesh and stucco
- Filing cabinets and metal equipment racks
- Transformers
- Heavy duty electric motors
- Fire walls and fire doors
- Concrete
- Refrigerators
- Sulphur plasma lighting (Fusion 2.4 GHz lighting systems)
- Air conditioning duct-work
- Other radio equipment
- Microwave ovens
- HVAC ducting
- Large transient elements such as forklifts or metal fabrications
- Other WLAN equipment

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

The site survey should also consider the three dimensional space occupied by the WLAN. For example, a multi-story building WLAN with different subnets per floor might require a different RF configuration than the same building with a single WLAN subnet per building. In the multiple subnet instances, a client attempting to roam to a different AP on the same floor might acquire an AP from an adjacent floor, which is not a desired behavior. Switching APs in a multi-subnet environment changes the roaming activity from a *seamless Layer 2 roam* to a *Layer 3 roam*, which in turn might disrupt sessions and require user intervention. This Layer 3 roaming in WLCs simplifies this deployment by allowing the client to maintain its IP address even if the AP to which it roams is on another subnet.

RF Deployment Best Practices

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

- The number of users versus throughput and a given AP. A common recommended number of users per AP is 15 to 25 for data-only users only and, for the 7920 VoIP (or similar voice devices) wireless handset, 7 to 8 voice users when data is present. This number should be used as a guideline and may vary depending on the handset in use. Check your handset requirements.

- The distance between APs can cause throughput variations for clients, based on the distance from the AP. The recommendation is to limit the AP data rate to the higher data rates to minimize data rate shifting.
- The number of APs depends on coverage and throughput requirements, which can vary. For example, Cisco System's internal information systems (IS) group currently uses six APs per 38,000 square feet of floor space for data-only operation.

**Note**

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

Manually Fine-Tuning WLAN Coverage

A number of factors can affect the WLAN coverage, as follows:

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

For a given data rate and location, the WLAN designer may alter power levels and/or elect to use a different antenna, to effect changes to the coverage area and/or coverage shape. Altering power levels or channel selection can be done manually as described below, or the Cisco Wireless Controller can do this automatically via the Radio Resource Management (RRM) algorithms, also referred to as Auto-RF. Cisco recommends the use of Radio Resource Management (RRM) to control the power level and channel, keeping in mind that the channel changing algorithm is highly dampened so that only a very disruptive (and persistent) interference source would cause a change to the channel topology, which in turn would cause clients to reassociate and any voice calls to be dropped. Changes in AP power do not impact clients. (See [Radio Resource Management \(Auto-RF\)](#), page 3-28 for more details).

Channel and Data Rate Selection

Channel selection depends on the frequencies that are permitted for a particular region. For example, the North American and ETSI 2.4 GHz channel sets permit allocation of three nonoverlapping channels: 1, 6, and 11 while the 5 GHz channel set permits 12 channels.

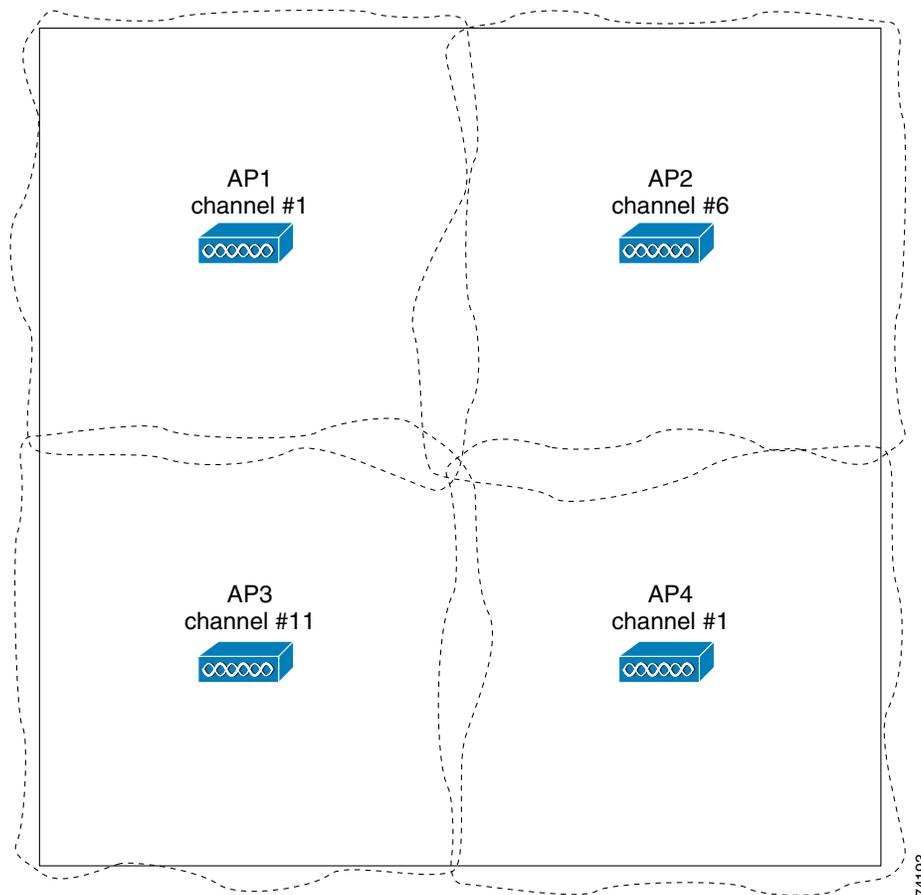
The channels should be allocated to the coverage cells as follows:

- Overlapping cells should use nonoverlapping channels.
- Where channels must be re-used in multiple cells, those cells should have minimal overlap with each other. [Figure 3-11](#) shows this pattern.

Recommendations for Channel Selection

Channel selection can be done manually, as described below.

Figure 3-11 Channel Allocated To APs



A site survey should be conducted using the same frequency plan as intended for the actual deployment. Some sites have high noise backgrounds which may prohibit the use of one or more channels. This provides a better estimate of how a particular channel at a particular location will react to the interference and the multipath. Channel selection also helps in planning for co-channel and the adjacent channel interference, and provides information about where you can reuse a frequency (see [Figure 3-12](#)).

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

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

Manual Channel Selection

Figure 3-12 shows a screenshot of the web page for configuring one of the 802.11b/g radios under the wireless selection. On the top right-hand side, channel 11 has been manually selected and the transmit power is set to 1, the highest level (8 sets the AP to the lowest level).



Note

The assignment method should normally be left at the global setting, unless there is a desire to manually control these settings. This allows the controller to dynamically change the channel number and transmit power as determined by the RRM. See [Radio Resource Management \(Auto-RF\), page 3-28](#) for more information.

Figure 3-12 Channel Assignment

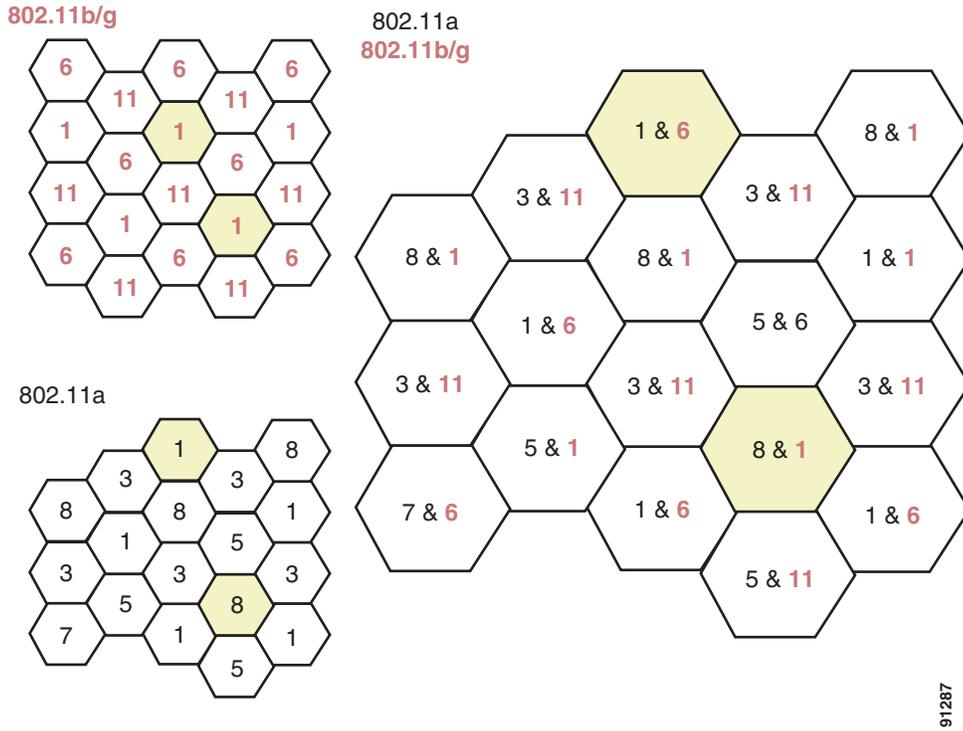
The screenshot displays the Cisco Systems web interface for configuring a 802.11b/g radio. The page is titled "802.11b/g Cisco APs > Configure" and includes a navigation menu with options like MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, and HELP. The left sidebar shows a tree view of configuration options under "Wireless", including Access Points, Bridging, Rogues, Clients, 802.11a, 802.11b/g, Country, and Timers. The main content area is divided into several sections:

- General:** AP Name (AP4_1ced.3294), Admin Status (Enable), Operational Status (UP), Site Config ID (0).
- Antenna:** Antenna Type (External), Diversity (Right), Antenna Gain (0 x 0.5 dBm).
- Management Frame Protection:** Version Supported (1), Protection Capability (AllFrames), Validation Capability (AllFrames).
- WLAN Override:** WLAN Override (Disable).
- RF Channel Assignment**:** Current Channel (11), Assignment Method (Custom).
- Tx Power Level Assignment:** Current Tx Power Level (1), Assignment Method (Custom).
- Performance Profile:** View and edit Performance Profile for this AP.

A note at the bottom of the page states: "** Note: Changing any of the parameters causes the Radio to be temporarily disabled and thus may result in loss of connectivity for some clients."

It is also possible to implement a dual-band deployment scheme, as shown in Figure 3-13. The top left portion of the diagram shows the 802.11b/g-only deployment, which uses the three nonoverlapping channels (channels 1, 6, and 11) to map out a pattern that has the least co-channel interference; that is, interference from an AP close by that is on the same channel, that is operating at sufficient power levels with its coverage pattern overlapping with that of another access point. It also shows an 802.11a deployment, which uses the eight nonoverlapping channels. The right side of the diagram illustrates how the channels would be mapped in a dual-band deployment.

Figure 3-13 Dual Band Deployment Diagram



Data Rate Selection

Figure 3-14 is a screenshot of the web page of the global 802.11b/g parameters. The data rate settings are shown on the right side of the screen.

Figure 3-14 Data Rate Assignment

The screenshot shows the Cisco Systems configuration page for 802.11b/g Global Parameters. The interface includes a navigation bar with tabs for MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, and HELP. The main content area is divided into sections: Wireless, Access Points, Bridging, Rogues, Clients, 802.11a, 802.11b/g, Country, and Timers. The 802.11b/g section is expanded to show Global Parameters. The General section includes settings for 802.11b/g Network Status (Enabled), 802.11g Support (Enabled), Beacon Period (100), DTIM Period (1), Short Preamble (Enabled), Fragmentation Threshold (2346), Pico Cell Mode (Disabled), and DTPC Support (Enabled). The CCX Location Measurement section has a Mode setting (Disabled). The Data Rates section lists rates from 1 Mbps to 54 Mbps with dropdown menus for Mandatory, Supported, or Disabled modes. A red note at the bottom explains the meaning of Mandatory and Supported modes.

802.11b/g Global Parameters

General

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

CCX Location Measurement

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

Data Rates**

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

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

153882

Mandatory, Supported, and Disabled Rate Modes

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

Data rates are expressed in megabits per second. You can set each data rate to one of three modes:

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

- **Supported**—Allows transmission at this rate for unicast packets only. The AP transmits only unicast packets at this rate; multicast and broadcast packets are transmitted at one of the data rates set to mandatory. The wireless clients always attempt to transmit and receive at the highest possible data rate. They negotiate with the AP for the highest data rate set to supported or mandatory to transmit and receive unicast packets. The wireless client devices are able to receive broadcast or multicast packets at any mandatory rate at or below the negotiated rate.
- **Disabled**—The AP does not transmit data at this rate.

Lowest and Highest Mandatory Rate Settings

Multiple clients associated to the AP can have completely different transmission rates, depending on interference, obstacles, or their distance from the AP. For example, if an 802.11b client is far from the AP and can only transmit and receive at a speed of 1 Mb/s because of the distance, it would be able to associate to the AP because the lowest mandatory rate (see [Figure 3-14](#)) is set to 1 Mb/s. If a second 802.11g client associates to the AP at 54 Mb/s, the AP would transmit broadcasts and multicasts at 1 Mb/s because this is the highest mandatory rate that all clients can receive. If the lowest mandatory rate was set to 5.5 Mb/s, the 802.11b client would not be able to associate to the AP because it could not receive broadcast packets at the lowest mandatory rate.

In [Figure 3-14](#), note that the highest mandatory setting is 11 Mb/s. The highest mandatory rate tells the AP what rate the client radios must be able to physically transmit at. This does not mean that they are actually transmitting and receiving packets at that rate, it just means that the radio physically supports that rate; the wireless client needs only to be able to receive packets at the lowest mandatory rate. 802.11b devices would be able to associate to the AP shown in [Figure 3-14](#) because their radios can physically transmit at 11 Mb/s. If a higher data rate (such as 18Mb/s) was set to mandatory, only 802.11g clients would be able to associate to the APs.

Setting any of the OFDM rates (rates above 11mb/s) to mandatory disables 802.11b connectivity. This can, for example, allow the administrator to exclude 802.11b clients from the AP by requiring an 802.11g data rate or setting a minimum transmission rate of all clients by disabling 802.11 rates. The reason this might be done is that the same 1500 byte packet at a lower data rate takes a longer time to transmit, and thus, lowers the effective data rate for all wireless clients associated to the AP.

Radio Resource Management (Auto-RF)

In the Cisco WLAN “split MAC” architecture (see [Chapter 2, “Cisco Unified Wireless Technology and Architecture,”](#)) the processing of 802.11 data and management protocols and access point capabilities is distributed between a lightweight access point and a centralized WLAN controller. More specifically, time-sensitive activities, such as probe response and MAC layer encryption, are handled at the access point. All other functions are sent to the controller, where system-wide visibility is required.

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



Note

An RF network group (or RF group) is not the same as a mobility group. A mobility group defines a mobility domain of 1–25 controllers in which a client would not be required to change IP address during a roaming event. This is accomplished by building Ethernet over IP tunnels for forwarding client data from an “anchor” controller to the “foreign” controller handling the new AP servicing the client.

Radio Resource Management (RRM), also known as Auto-RF, can adjust the channel (dynamic channel assignment) and power (dynamic transmit power control) to maintain the RF coverage area. It adjusts the power level of the AP to maintain a baseline signal strength with neighboring APs at -65 dBm (configurable) (See [Overview of Auto-RF Operation, page 3-29](#)). It adjusts the channel of the AP when it notices nearby interference sources on the channel on which the AP is currently located. It continues to optimize the RF coverage for the best reception and throughput for the wireless network.

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

For more detailed information on Radio Resource Management (Auto-RF), see the following URL: http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a008072c759.shtml.

Overview of Auto-RF Operation

Each controller is configured with an RF network group name (called RF Network Name under the WLC Controller -> General menu). In each RF group (if Group Mode is enabled), the controllers elect a leader and form an RF domain. The function of the leader is to collect the network-wide neighbor information from a group of controllers and do the channel/power computation for an optimal system-wide map. If Group Mode is not enabled, the controllers run computations based only on the neighbor data gathered from the APs connected via LWAPP, trying to optimize the signal to -65 dBm between APs.

The APs transmit Radio Resource Management (RRM) neighbor packets at full power at regular intervals. These messages contain a field that is a hash of the RF group name, BSSID, and time stamp. The APs accept only RRM neighbor packets sent with this RF network name.

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

[Table 3-7](#) provides a summary of the various functions of the devices in the system.



Note

TPC performs only downward power level adjustments. Coverage hole detection and correction increases power levels on APs. If a client is associated at a low RSSI level for a significant period of time (and thus has not roamed), it is assumed they cannot find another AP and must be in a coverage hole, and power is increased to help the client communicate.

Auto-RF should not be confused with Rogue Detection (channel scanning), which is done separately from the auto-RF algorithm. APs perform rogue detection by periodically monitoring all country-specific channels (channel scanning). The APs go “off-channel” for a period not greater than 60 ms to listen to the other channels. Packet headers collected during this time are sent to the controller, where they are analyzed to detect rogue access points, whether service set identifiers (SSIDs) are broadcast or not, rogue clients, ad-hoc clients, and interfering access points.

By default, each access point spends approximately 0.2 percent of its time off-channel. This is statistically distributed across all access points so that no two adjacent access points are scanning at the same time, which can adversely affect WLAN performance. Packets received by the AP from clients are

forwarded to the controller with the LWAPP status field filled in, which provides the controller with radio information including RSSI and signal-to-noise ratio (SNR) for all packets received by the AP during reception of the packet.

Table 3-7 Device Function

Device	Functions
RF Group Leader	Collects data from WLCs in the RF group and analyzes it for TX Power Control (TPC) and Dynamic Channel Assignment (DCA) system-wide. TPC adjusts power levels only downward.
Local WLC	Collects data and runs the Coverage Hole Detection and Correction algorithm. Adjusts power levels upward if necessary for clients
Light-weight access point	<ul style="list-style-type: none"> • Sends neighbor messages on all channels at full power at configured interval • Verifies neighbor hash on received neighbor messages • Scans configured channels for noise, interference, and IDS/rogue detection and alerts if profile fails

Auto-RF Variables and Settings

Auto-RF can be turned on and off via the global setting on the Channel Selection (**Wireless > 802.11b > Configure**) web page (see [Figure 3-12](#)). You can manually set the channel and transmit level for the AP from this web page. Additionally, it can be turned off and on from the global Auto-RF web page. Remember that Auto-RF is per band and RF group computations are done for both the 802.11b/g band and another set of computations for 802.11a. The two radios do not have to share to have the same configuration. But these configurations are applied to every AP associated to the controller. Auto-RF configuration variables are shown on the global parameters Auto-RF configuration page (see [Figure 3-15](#)).

The first set of variables on the Auto-RF configuration web page corresponds to the RF group. These determine whether the controller joins the dynamic grouping with the other controllers. The dynamic grouping helps the controller find out about APs that are neighbors but might be associated to another controller in the mobility group. If this is disabled, the controller only optimizes the parameters of the access points that it knows about (that is, the ones that are associated to it). The group leader indicates the MAC address of the elected leader. You can find the MAC address of the controller on the inventory web page (you can reach the web page by clicking on **Controller** at the top menu and then **Inventory**).

The Auto-RF configuration web page is divided into three pages, or sections, with a scroll bar that is used to move among the three pages. The first page (see [Figure 3-15](#)) is for dynamic channel assignment. This allows the controller to automatically change the channel that the AP is on (for more information, see [Dynamic Channel Assignment, page 3-33](#)).

Figure 3-15 Auto-RF (Page 1)

The screenshot shows the Cisco Auto-RF configuration interface. The main heading is "802.11b/g Global Parameters > Auto RF". On the left, there is a navigation menu with categories: Wireless, Access Points, Bridging, Rogues, Clients, 802.11a, 802.11b/g, Country, and Timers. The main content area is divided into sections: "RF Group" and "RF Channel Assignment".

RF Group

Group Mode	<input checked="" type="checkbox"/> Enabled
Group Update Interval	600 secs
Group Leader	00:0b:85:40:40:00
Is this Controller a Group Leader	Yes
Last Group Update	363 secs ago

RF Channel Assignment

Channel Assignment Method	<input checked="" type="radio"/> Automatic Every 600 sec <input type="radio"/> On Demand Invoke Channel Update now <input type="radio"/> OFF
Avoid Foreign AP interference	<input checked="" type="checkbox"/> Enabled
Avoid Cisco AP load	<input type="checkbox"/> Enabled
Avoid non-802.11b noise	<input checked="" type="checkbox"/> Enabled
Signal Strength Contribution	Enabled
Channel Assignment Leader	00:0b:85:40:40:00
Last Channel Assignment	363 secs ago

On the right side, there are links for "RF Group Members" and "MAC Address" (00:0b:85:40:40:00). At the bottom, there is a link for "Tx Power Level Assignment".

Following the RF channel assignment is the section for assigning the transmit (tx) power level (see Figure 3-16). On this web page, the power level can be fixed for all APs, or it can be automatically adjusted. The web page also indicates the number of neighbors the AP has and the power thresholds for which it is adjusting.

Figure 3-16 Auto-RF (Page 2)

The screenshot shows the "Tx Power Level Assignment" configuration page. The main heading is "Tx Power Level Assignment". On the left, the navigation menu is similar to the previous page, but the "Clients" section is expanded to show "802.11a" and "802.11b/g" sub-sections.

Tx Power Level Assignment

Power Level Assignment Method	<input checked="" type="radio"/> Automatic Every 600 sec <input type="radio"/> On Demand Invoke Power Update now <input type="radio"/> Fixed <input type="text" value="1"/>
Power Threshold	-65 dBm
Power Neighbor Count	3
Power Update Contribution	SNI.
Power Assignment Leader	00:0b:85:40:40:00
Last Power Level Assignment	363 secs ago

Profile Thresholds

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

The third web page is for profile thresholds. The controller analyzes the information passed to it by the APs and determines a pass or fail status for each of these thresholds. These pass/fail profiles are best seen in the output of the **show ap auto-rf radio ap_name** command (see the following sample). The same information can be seen in graphical form on the **Monitor > 802.11b/g Radios > Detail** web page.

Sample show ap auto-rf Command Output

```

show>ap auto-rf 802.11b <access point name>
Number of Slots . . . . . 2
AP Name . . . . . <AP name>
MAC Address . . . . . 00:0b:85:1b:df:c0
Radio Type . . . . . RADIO_TYPE_80211b/g
Noise Information
  Noise Profile . . . . . PASSED
  Channel 1 . . . . . -93 dBm
  Channel 2 . . . . . -90 dBm
  .
  .
  Channel 11 . . . . . -95 dBm
Interference Information
  Interference Profile . . . . . FAILED
  Channel 1 . . . . . -69 dBm @ 31 % busy
  Channel 2 . . . . . -58 dBm @ 26 % busy
  .
  .
  Channel 11. . . . . -68 dBm @ 26 % busy
Load Information
  Load Profile . . . . . PASSED
  Receive Utilization . . . . . 0 %
  Transmit Utilization . . . . . 0 %
  Channel Utilization . . . . . 26 %
  Attached Clients . . . . . 2 clients
Coverage Information
  Coverage Profile . . . . . PASSED
  Failed Clients . . . . . 0 clients
Client Signal Strengths
  RSSI -100 dBm. . . . . 0 clients
  RSSI -92 dBm . . . . . 0 clients
  .
  .
  RSSI -52 dBm . . . . . 1 clients
Client Signal To Noise Ratios
  SNR 0 dBm . . . . . 0 clients
  SNR 5 dBm . . . . . 0 clients
  SNR 10 dBm . . . . . 0 clients
  .
  .
  SNR 45 dBm . . . . . 1 clients
Nearby APs
Radar Information
Channel Assignment Information
  Current Channel Average Energy . . . . . -68 dBm
  Previous Channel Average Energy . . . . . -51 dBm
  Channel Change Count . . . . . 21
  Last Channel Change Time . . . . . Thu Mar 9 12:18:03 2006
  Recommend Best Channel . . . . . 11
RF Parameter Recommendations
  Power Level . . . . . 1

```

RTS/CTS Threshold	2347
Fragmentation Threshold	2346
Antenna Pattern	0

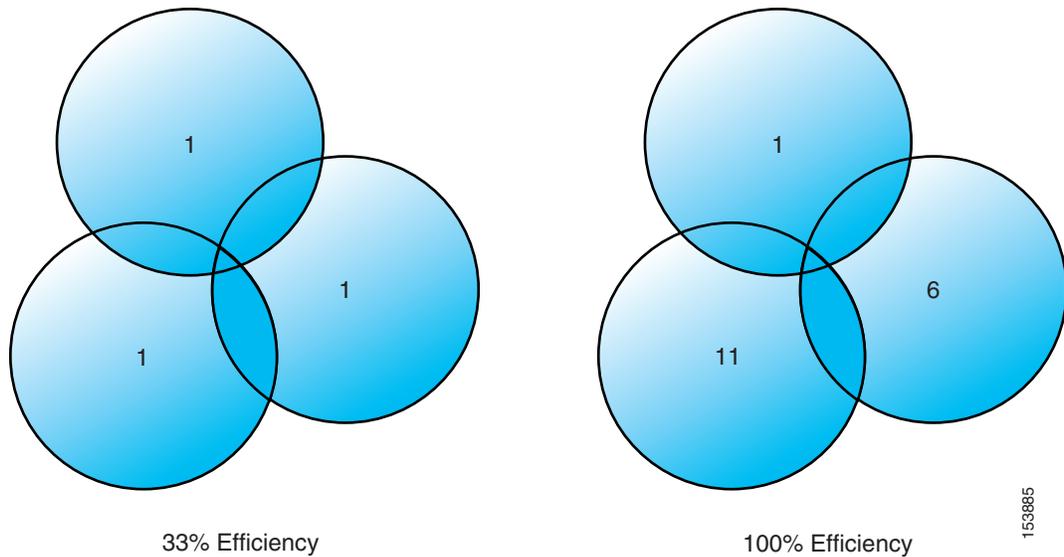
The following sections describe some of the Auto-RF variables.

Dynamic Channel Assignment

802.11 MAC uses Carrier-Sense Multiple Access/Collision Avoidance (CSMA/CA). With CSMA/CA, two access points on the same channel (in the same vicinity) get half the capacity of two access points on different channels because of the shared wireless channel. This becomes an issue, for example, when someone reading an e-mail in one business affects the performance of an access point in a neighboring business. Even though these are completely separate networks, someone sending traffic on the common channel can cause network busy signals to clients in the nearby business. This is not to be confused with beacon traffic, which commonly extends well beyond cell range and/or collision domain range. The controllers address this problem and other co-channel interference issues by dynamically allocating access point channel assignments to avoid conflict. Because the controller, or a designated controller called an RF Group Leader, has a system-wide visibility, channels are “reused” to avoid wasting scarce RF resources. In other words, the access point is allocated a different channel, far from the channel used by the neighboring business.

The dynamic channel assignment capabilities of the controller are also useful in minimizing co-channel interference between adjunct access points. For example, with 802.11g, nearby APs cannot both simultaneously use the same channel and receive at 54 Mbps because of interference with each other. By assigning channels, the controller keeps adjacent channels separated, avoiding this problem, as shown in Figure 3-17.

Figure 3-17 Dynamic Channel Assignment



The controller examines a variety of real-time RF characteristics to efficiently handle channel assignments. These include:

- **Noise**—This limits signal quality at the client and access point, if received in an amount greater than approximately -65 dBm, though this can vary quite easily with directional antennas, range, and periodicity. There are numerous types and effects of interference. An increase in noise reduces the effective cell size at the intervals determined by the controller when it reassesses a radiating

environment. By optimizing channels to avoid noise sources, the controller can optimize coverage while maintaining system capacity. If a channel is unusable because of excessive noise, that channel can be avoided. If other wireless networks are present, the controller shifts the usage of channels to complement the other networks. For example, if one network is on Channel 6, an adjacent WLAN is assigned Channel 1 or 11. This increases the capacity of the network by limiting the sharing of frequencies. If a channel is used so much that no capacity is available, the Cisco Wireless LAN Controller might choose to avoid this channel.

- **Client load**—Client load is taken into account when changing the channel structure to minimize the impact on the clients currently on the WLAN system. The controller periodically monitors the channel assignment in search of the best assignments. Change occurs only if it significantly improves the performance of the network or corrects the performance of a poorly performing access point.

The controller combines the RF characteristic information to make system-wide decisions. The end result is an optimal channel configuration in a three-dimensional space, where access points on the floor above and below play a major factor in an overall WLAN configuration.

Interference Detection and Avoidance

Interference is defined as any 802.11 traffic that is not part of the Cisco WLAN system; this includes a rogue access point, or a neighboring WLAN. It can also include non-802.11 sources such as certain microwave ovens or cordless phones. It can in certain instances also include various sources of electromagnetic interference (EMI) such as arc welders or federal/military radar facilities. APs are constantly scanning all channels looking for major sources of interference.

If the amount of 802.11 interference hits a predefined threshold, the controller attempts to rearrange channel assignments to optimize system performance in the presence of the interference. This might result in adjacent APs being on the same channel, but logically this is a better choice than staying on a channel that is totally unusable because of an interfering access point.

The controller can respond to a rogue AP on channel 11 by shifting the closest APs to channel 1 or channel 6.

Dynamic Transmit Power Control

The correct AP power settings are essential to maintaining the coverage area, not only to ensure correct (not maximum) amount of power covering an area, but also to ensure that an excess of power is not used, which usually adds significant amounts of noise to the radiating area. AP power settings also control network redundancy by helping to ensure real-time failover in the event of the loss of an AP. The controller is used to dynamically control the AP transmit power level based on real-time WLAN conditions. In normal instances, power can be kept low to gain extra capacity and reduce interference among the APs. RRM attempts to balance access points such that they see their neighbors at -65 dBm. If a failed access point is detected, power can be automatically increased on surrounding access points to fill the coverage gap created by the loss of the AP.

RRM algorithms are designed to create the optimal user experience. For example, if the power of an access point is turned down to Level 4 (where Level 1 = highest and Level 8 = lowest) and the received signal strength indicator (RSSI) value or a user drops below an acceptable threshold, the access point power is increased to provide a better experience to that client.

The change in power settings can be communicated to the clients (via CCX) so that clients see a balanced cell because of the matched transmit power setting and data rate settings between the AP and client. When a client device associates to the access point, the access point sends the maximum power level

setting to the client. If the access points adjust for high data rates and low transmit power but similar adjustments are not made on the clients, the cell might be unbalanced because the client signal can be greater than the access point signal and cause undue interference with neighboring cells.

**Note**

If the access point is configured to disable data rates 1, 2, and 5.5, clients are not required to transmit only at 11 Mbps. In fact, older client devices might not be programmed to recognize the access point configuration for data rates.

Coverage Hole Detection and Correction

If clients on an access point are detected at low RSSI levels, this indicates the existence of an area where clients are continually getting poor signal coverage, without having a viable location to which to roam. The controller adjusts access point power levels to correct the detected hole.

Client and Network Load Balancing

The IEEE standard did not define the process or reasons for client roaming, and therefore it cannot be easily predicted what clients will do in any given situation. For example, all users in a conference room can associate with a single access point because of its close proximity, ignoring other access points that are farther away but much less used.

The controller has a centralized view of client distribution across all access points. This can be used to influence where new clients attach to the network if there are multiple “good” APs available. If configured, the controller can proactively ‘herd’ clients to new access points to improve WLAN performance. This results in a smooth distribution of capacity across an entire wireless network. Keep in mind that this load balancing is done at client association, not while a client is connected.

