Chapter 17: Wireless Signals and Modulation
Chapter 18: Wireless Infrastructure
Chapter 19: Understanding Wireless Roam and Location Services

Instructor Materials

CCNP Enterprise: Core Networking

Understanding Basic Wireless Theory

- Wireless signals travel as electromagnetic waves through the air from sender to receiver.
- Frequency is a fundamental property of the waves involved in a wireless link.

Basic Wireless Concepts

In RF wireless communications, the sender (a transmitter) sends an alternating current into a section of wire (an antenna), which sets up moving electric and magnetic fields that propagate out and away from the antenna as traveling waves.

- The electric and magnetic fields travel along together and are always at right angles to each other, as shown in *Figure 17-3*. The signal must keep changing, or alternating, by cycling up and down, to keep the electric and magnetic fields cycling and pushing ever outward.
- Electromagnetic waves do not travel strictly in a straight line. Instead, they travel by expanding in all directions away from the antenna.
- In free space, the electromagnetic waves expand outward in all three dimensions.

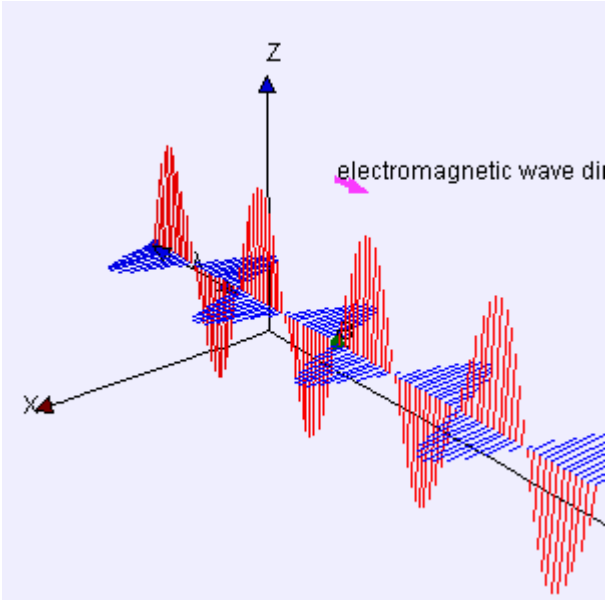

Figure 17-3 *Traveling Electric and Magnetic Waves*

Basic Wireless Concepts (Cont.)

The waves produced from a tiny point antenna expand outward in a spherical shape. The waves will eventually reach the receiver, in addition to many other locations in other directions.

- Figure 17-4 shows a simple idealistic antenna that is a single point, which is connected at the end of a wire at the sender.
- At the receiving end of a wireless link, the process is reversed. As the electromagnetic waves reach the receiver's antenna, they induce an electrical signal. If everything works right, the received signal will be a reasonable copy of the original transmitted signal.

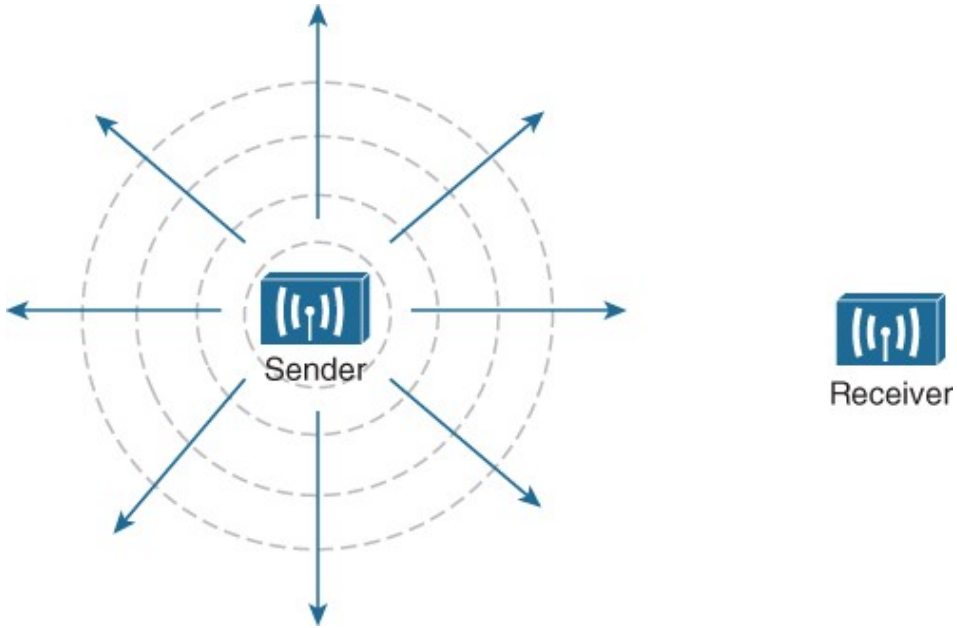

Figure 17-4 *Wave Propagation with an Idealistic Antenna*

Understanding Frequency

The waves involved in a wireless link can be measured and described in several ways. One fundamental property is the frequency of the wave, or the number of times the signal makes one complete up and down cycle in 1 second.

- A cycle can begin as the signal rises from the center line, falls through the center line, and rises again to meet the center line.
- A hertz (Hz) is the most commonly used frequency unit and corresponds to the number of cycles per second.
- In Figure 17-5, suppose that 1 second has elapsed, as shown. During that 1 second, the signal progressed through four complete cycles. Therefore, its frequency is 4 cycles/second, or 4 hertz.

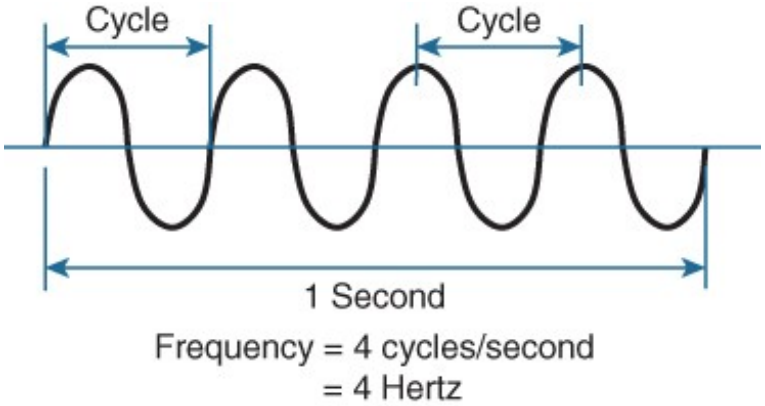

Figure 17-5 *Cycles Within a Wave*

Frequency Unit Names

Frequency can vary over a very wide range. As frequency increases by orders of magnitude, the numbers can become quite large. To keep things simple, the frequency unit name can be modified to denote an increasing number of zeros, as listed in Table 17-2.

Table 17-2 Frequency Unit Names

Unit	Abbreviation	Meaning
Hertz	Hz	Cycles per second
Kilohertz	kHz	1000 Hz
Megahertz	MHz	1,000,000 Hz
Gigahertz	GHz	1,000,000,000 Hz

Continuous Frequency Spectrum

Frequency (Hz)	Frequency Notation
10^{22}	
10^{21}	
10^{20}	
10^{19}	
10^{18}	
10^{17}	
10^{16}	
10^{15}	
10^{14}	
10^{13}	
10^{12}	100 GHz
10^{10}	10 GHz
10^{9}	1 GHz
10^{8}	100 MHz
10^{7}	10 MHz
10^{6}	1 MHz
10^{5}	100 kHz
10^{4}	10 kHz
10^{3}	1 kHz
10^{2}	100 Hz
10^{1}	10 Hz
0	0 Hz

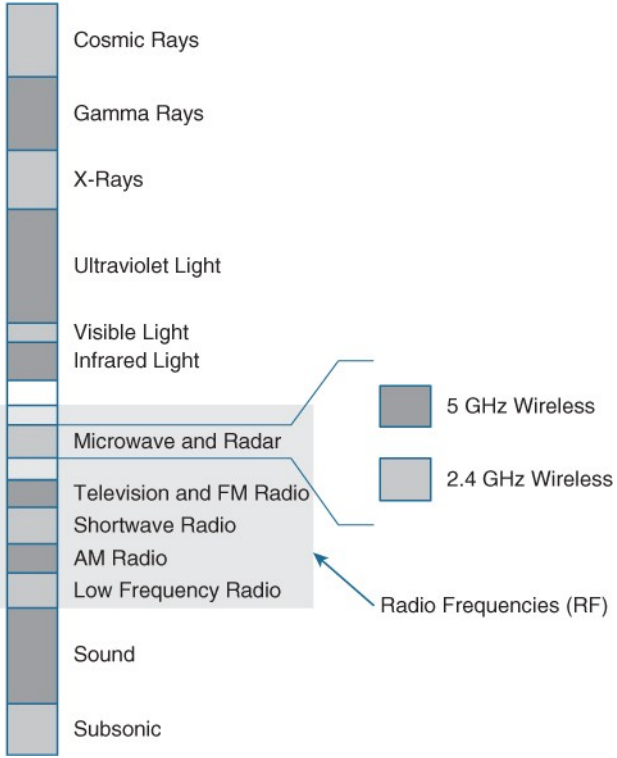

Figure 17-6 *Continuous Frequency Spectrum*

Figure 17-6 shows a simple representation of the continuous frequency spectrum ranging from 0 Hz to 10^{22} (or 1 followed by 22 zeros) Hz. At the low end of the spectrum are frequencies that are too low to be heard by the human ear, followed by audible sounds. The highest range of frequencies contains light, followed by X, gamma, and cosmic rays.

The frequency range from around 3 kHz to 300 GHz is commonly called radio frequency (RF). It includes many different types of radio communication, such as low-frequency radio, AM radio, shortwave radio, television, FM radio, microwave, and radar.

Frequency Bands for Wireless LANs

One of the two main frequency ranges used for wireless LAN communication lies between 2.400 and 2.4835 GHz. This is usually called the 2.4 GHz band, even though it does not encompass the entire range between 2.4 and 2.5 GHz.

The other wireless LAN range is usually called the 5 GHz band because it lies between 5.150 and 5.825 GHz. The 5 GHz band actually contains the following four separate and distinct bands:

- 5.150 to 5.250 GHz
- 5.250 to 5.350 GHz
- 5.470 to 5.725 GHz
- 5.725 to 5.825 GHz

Most of the 5 GHz bands are contiguous except for a gap between 5.350 and 5.470. At the time of this writing, this gap exists and cannot be used for wireless LANs.

Understanding Frequency - Channels

Frequency bands are usually divided up into a number of distinct channels. Each channel is known by a channel number and is assigned to a specific frequency. As long as the channels are defined by a national or international standards body, they can be used consistently in all locations.

Figure 17-7 shows the channel assignment for the 2.4 GHz band that is used for wireless LAN communication. The band contains 14 channels numbered 1 through 14, each assigned a specific frequency.

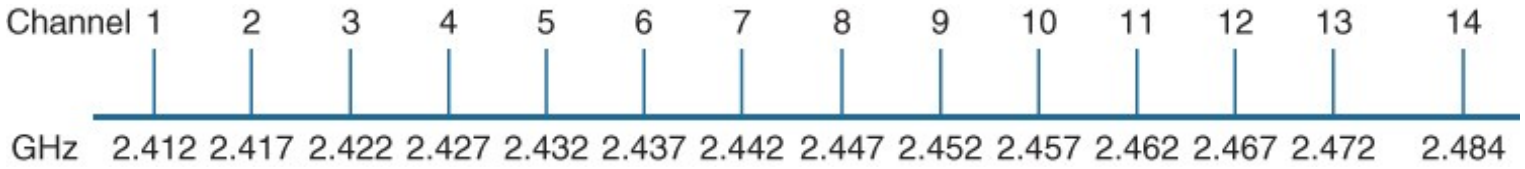

Figure 17-7 *Example of Channel Spacing in the 2.4 GHz Band*

Understanding Frequency - Bandwidth

The actual frequency range needed for the transmitted signal is known as the signal *bandwidth*, as shown in Figure 17-8. As its name implies, bandwidth refers to the width of frequency space required within the band.

In wireless LANs, the signal bandwidth is defined as part of a standard. Even though the signal might extend farther above and below the center frequency than the bandwidth allows, wireless devices will use something called a spectral mask to ignore parts of the signal that fall outside the bandwidth boundaries.

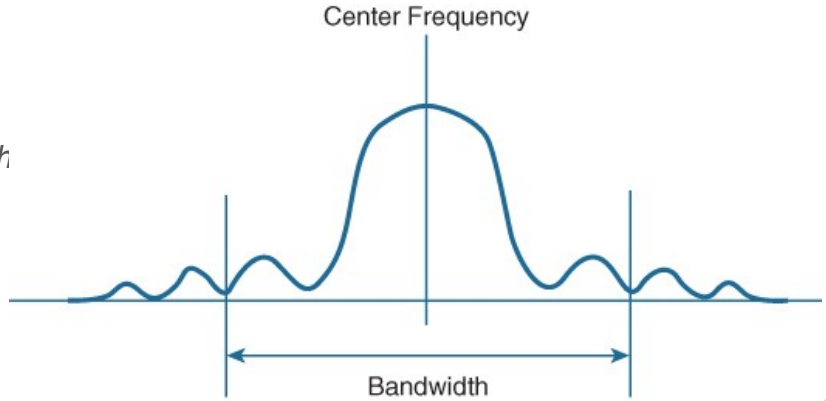

Figure 17-8 *Signal Bandwidth*

Understanding Frequency – Overlapping Channels

Ideally, the signal bandwidth should be less than the channel width so that a different signal could be transmitted on every possible channel, with no chance that two signals could overlap and interfere with each other.

When the signal bandwidth is wider than the channel assignment, the signals overlap each other, as shown in Figure 17-10. Because of this, signals on adjacent channels cannot possibly coexist without interfering with each other.

Signals must be placed on more distant channels to prevent overlapping, thus limiting the number of usable channels in the band.

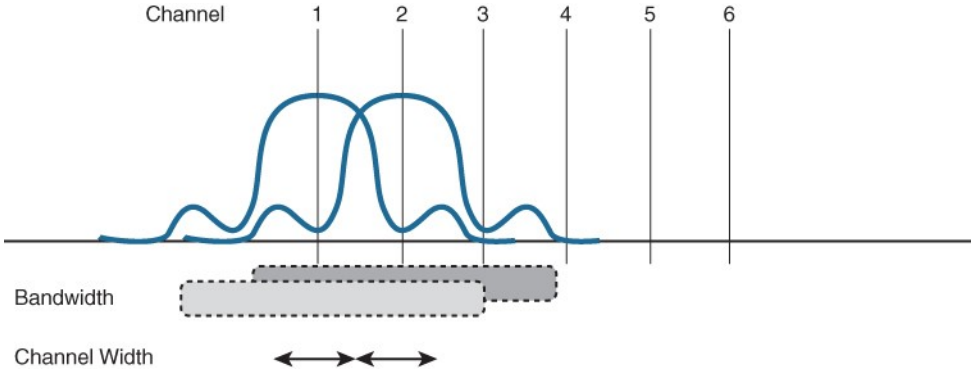

Figure 17-10 *Overlapping Channel Spacing*

Understanding Phase

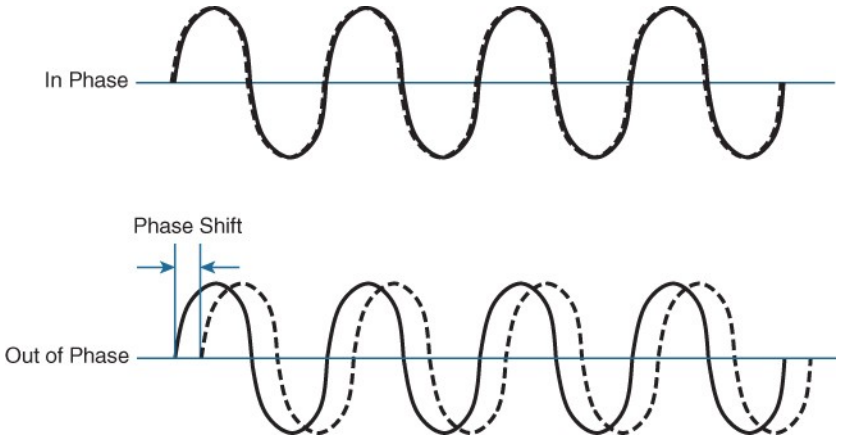

Figure 17-11 *Signals In and Out of Phase*

RF signals are very dependent upon timing because they are always in motion.

The phase of a signal is a measure of shift in time relative to the start of a cycle.

When two identical signals are produced at exactly the same time, their cycles match up and they are said to be in phase with each other. If one signal is delayed from the other, the two signals are said to be out of phase.

Signals that are in phase tend to add together, whereas signals that are 180 degrees out of phase tend to cancel each other out.

Measuring Wavelength

Wavelength is a measure of the physical distance that a wave travels over one complete cycle. Wavelength is usually designated by the Greek symbol lambda (λ).

Regardless of the frequency, RF waves travel at a constant speed. In a vacuum, radio waves travel at exactly the speed of light; in air, the velocity is slightly less than the speed of light.

Wavelength decreases as the frequency increases. As the wave cycles get smaller, they cover less distance.

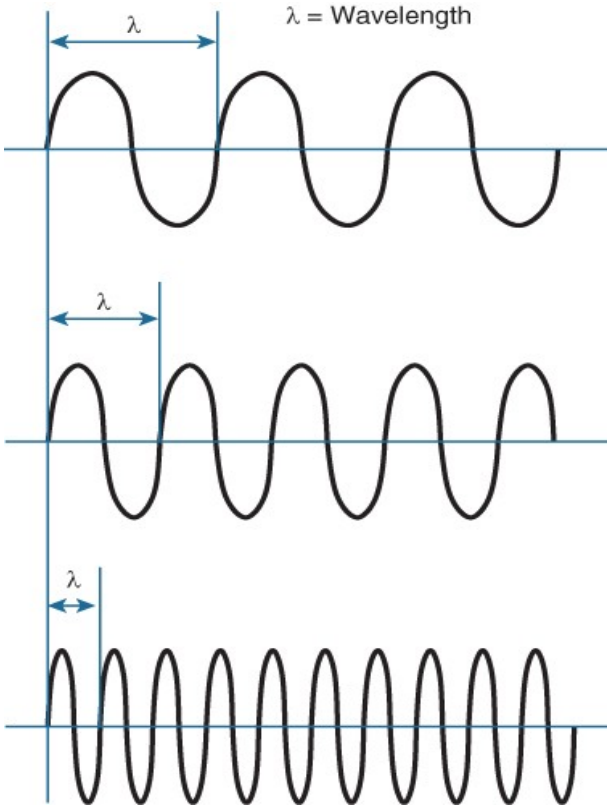

Figure 17-12 *Examples of Increasing Frequency and Decreasing Wavelength*

Understanding RF Power and dB

Use of the dB logarithmic functions:

- The strength of a wave can be measured as its amplitude, the top to bottom peak.
- The strength of an RF signal is usually measured by its power, in watts (W).
- When power is measured in watts (W) or milliwatts (mW), it is considered to be an absolute power measurement.
- Because absolute power values can fall anywhere within a huge range, from a tiny decimal number to hundreds, thousands, or greater values, we use logarithmic functions to transform exponential ranges into linear ones.
- The decibel (dB) is a handy function that uses logarithms to compare one absolute measurement to another. It was originally developed to compare sound intensity levels, but it applies directly to power levels, too.

Understanding RF Power and dB (Cont.)

The following equation is used to calculate a dB value, where P1 and P2 are the absolute power levels of two sources:

$$dB = 10(\log P_2 - \log P_1)$$

P2 represents the source of interest, and P1 is usually called the reference value or the source of comparison. The difference between the two logarithmic functions can be rewritten as a single logarithm of P2 divided by P1, as follows:

$$dB = 10.\log\left(\frac{P_2}{P_1}\right)$$

The ratio of the two absolute power values is computed first; then the result is converted onto a logarithmic scale. The ratio or division form of the equation is the most commonly used in the wireless engineering world.

Understanding RF Power and dB (Cont.)

Important dB laws are as follows:

Law of Zero: A value of 0 dB means that the two absolute power values are equal. If the two power values are equal, the ratio inside the logarithm is 1, and the log10(1) is 0. This law is intuitive; if two power levels are the same, one is 0 dB greater than the other.

Law of 3s: A value of 3 dB means that the power value of interest is double the reference value; a value of −3 dB means the power value of interest is half the reference. When P2 is twice P1, the ratio is always 2. Therefore, 10log10(2) = 3 dB. When the ratio is 1/2, 10log10(1/2) = −3 dB.The Law of 3s is not very intuitive, but is still easy to learn. Whenever a power level doubles, it increases by 3 dB. Whenever it is cut in half, it decreases by 3 dB.

Law of 10s: A value of 10 dB means that the power value of interest is 10 times the reference value; a value of −10 dB means the power value of interest is 1/10 of the reference.

- When P2 is 10 times P1, the ratio is always 10. Therefore, 10log10(10) = 10 dB.
- When P2 is one tenth of P1, then the ratio is 1/10 and 10log10(1/10) = −10 dB.
- The Law of 10s is intuitive because multiplying or dividing by 10 adds or subtracts 10 dB, respectively.

Understanding RF Power and dB (Cont.)

When absolute power values multiply, the dB value is positive and can be added. When the power values divide, the dB value is negative and can be subtracted. Table 17-3 summarizes the useful dB comparisons.

Table 17-3 Power Changes and Their Corresponding dB Values

Power Change	dB Value
=	0 dB
× 2	+3 dB
/ 2	−3 dB
× 10	+10 dB
/ 10	−10 dB

Measuring Power Changes Along the Signal Path

A transmitter, its antenna, and the cable that connects them are all discrete components that not only propagate an RF signal but also affect its absolute power level.

- Transmitter power is usually a known value, expressed in mW.
- Antenna connected to a transmitter provides some amount of gain (it is measured by comparing its performance with that of a reference antenna (usually an *isotropic* antenna), then computing a value in dB.
- Some signal loss occurs due to the physical qualities of the cable that connects an antenna to a transmitter.
- Antenna connected to a receiver has also some gain similarly like transmitter one.
- Some signal loss occurs due to the physical qualities of the cable that connects an antenna to a receiver.

Measuring Power Changes Along the Signal Path

Radio link budget formula *(energetická bilance radiového spoje)*

Received power (dBm)=Transmitted power (dBm)+Gains (dB)−Losses (dB)

$$P_{RX}=P_{TX}+G_{TX}+G_{RX}-L_{TX}-L_{FS}-L_P-L_{RX}$$

Where:

P_{RX} = received power (dBm)
P_{TX} = transmitter output power (dBm)
G_{TX} = transmitter antenna gain (dBi)
G_{RX} = receiver antenna gain (dBi)
L_{TX} = transmit feeder and associated losses (feeder, connectors, etc.) (dB)
L_{FS} = free space loss or path loss (dB)
L_P = miscellaneous signal propagation losses (these include fading margin, polarization mismatch, losses associated with medium through which signal is travelling, other losses...) (dB)
L_{RX} = receiver feeder and associated losses (feeder, connectors, etc.) (d)B
r = path length (distance between the antennas)
λ = wavelength

where $L_{FS}=\left(\frac{4\pi r}{\lambda}\right)^2$

Understanding RF Power and dB

Beyond comparing two transmitting sources, a network engineer must be concerned about the RF signal propagating from a transmitter to a receiver. The dB formula to compare the received signal strength to the transmitted signal strength is:

$$dB = 10\log_{10}\left(\frac{0.000031623\,mW}{100\,mW}\right) = -65\,dB$$

The absolute power values at the transmitter and receiver can be converted to dBm, the results of which are shown in Figure 17-19. Notice that the dBm values can be added along the path: The transmitter dBm plus the net loss in dB equals the received signal in dBm.

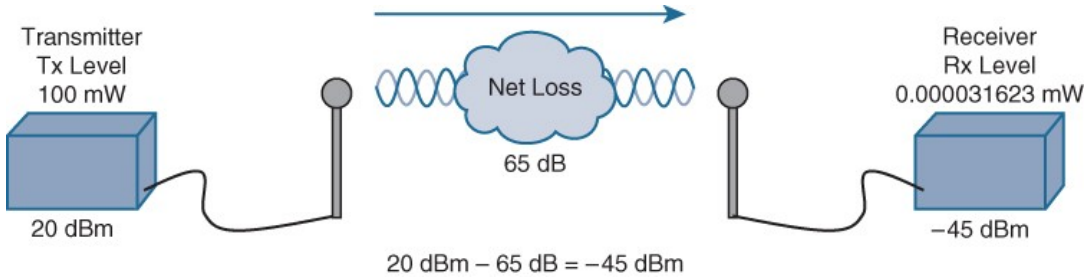

Figure 17-19 *Subtracting dB to Represent a Loss in Signal Strength*

Measuring Power Changes Along the Signal Path (Cont.)

The effective isotropic radiated power (EIRP) is the actual power level that will be radiated from the antenna. This value is calculated as a combination of transmitter power, the loss from the length of cable, and the antenna gain. The formula to calculate EIRP: EIRP = Tx Power – Tx Cable + Tx Antenna. EIRP is regulated by government agencies in most countries, a signal cannot exceed the maximum allowable EIRP.

A link budget is the power levels across the entire path from transmitter to receiver. To calculate the received signal strength, begin with the transmitter power expressed in dBm, add or subtract the dB components along the signal path to find the signal strength that arrives at the receiver.

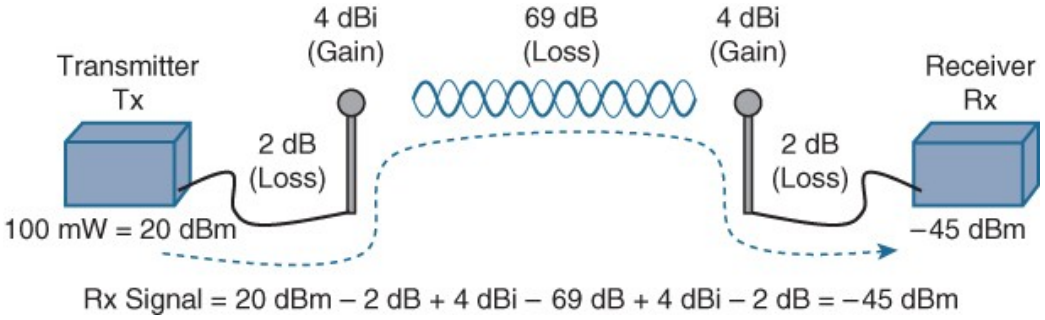

Figure 17-22 *Example of Calculating Received Signal Strength*

Free Space Path Loss

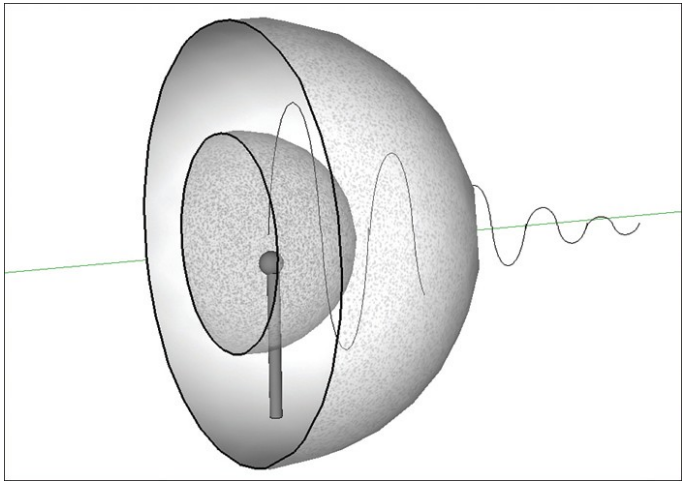

Figure 17-23 *Free Space Loss Due to Wave Spreading*

- Whenever an RF signal is transmitted from an antenna, its amplitude decreases as it travels through free space.
- Even if there are no obstacles in the path between the transmitter and receiver, the signal strength will weaken. This is known as free space path loss.
- If the antenna is a point the RF data wave energy travels in every direction from the antenna. The wave that is produced takes the form of a sphere.
- As energy is transmitted from the antenna, the sphere expands in free space.
- Regardless of the antenna used, the amount of free space path loss through free space is consistent. Two facts:
  - Free space path loss is an exponential function; the signal strength falls off quickly near the transmitter but more slowly farther away.
  - The loss is a function of distance and frequency only.

Free Space Path Loss (Cont.)

Free space path loss is greater in the 5 GHz band than it is in the 2.4 GHz band. In the equation, as the frequency increases, so does the loss in dB.

Figure 17-24 shows the range difference, where both transmitters have an equal EIRP. The dashed circles show where the effective range ends, at the point where the signal strength of each transmitter is equal.

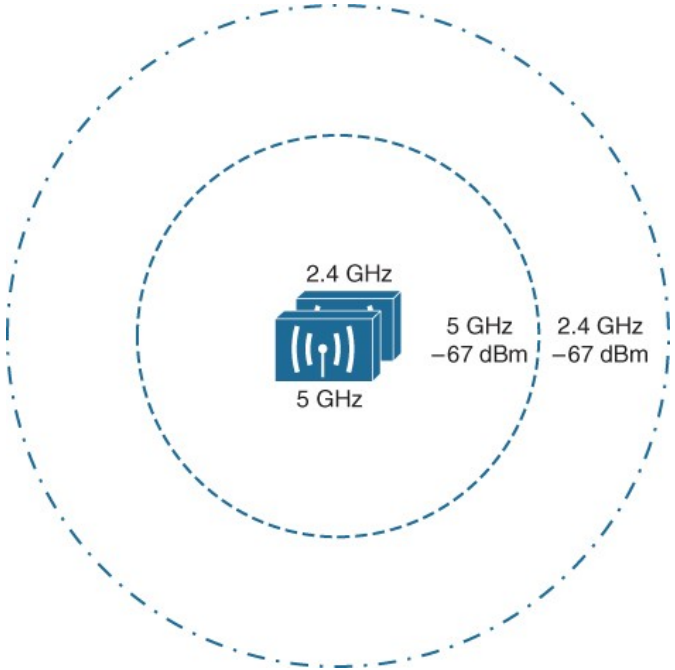

Figure 17-24 *Effective Range of 2.4 GHz and 5 GHz Transmitters*

Power Levels at the Receiver

Receivers usually measure a signal's power level according to the received signal strength indicator (RSSI) scale.

The RSSI value is defined in the 802.11 standard as an internal 1-byte relative value ranging from 0 to 255, where 0 is the weakest and 255 is the strongest. The range of RSSI values can vary between one hardware manufacturer and another.

Every receiver has a sensitivity level, or a threshold that divides intelligible, useful signals from unintelligible ones. As long as a signal is received with a power level that is greater than the sensitivity level, chances are that the data from the signal can be understood correctly.

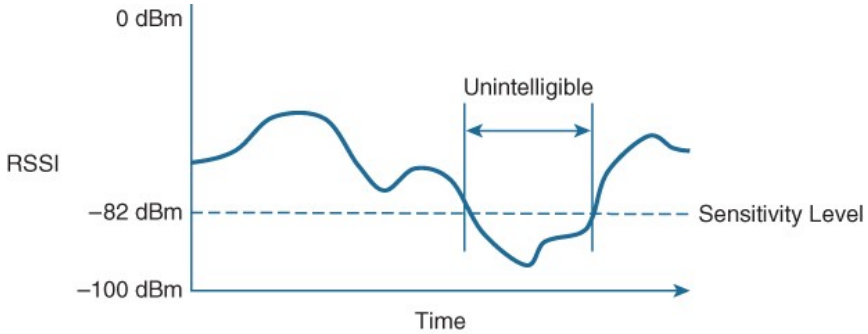

Figure 17-25 *Example of Receiver Sensitivity Level*

Power Levels at the Receiver (Cont.)

The RSSI value focuses on the expected signal alone, without regard to any other signals that may also be received. All other signals that are received on the same frequency as the one you are trying to receive are simply viewed as noise. The noise level, or the average signal strength of the noise, is called the noise floor.

The difference between the signal and the noise is called the signal-to-noise ratio (SNR), measured in dB. A higher SNR value is preferred.

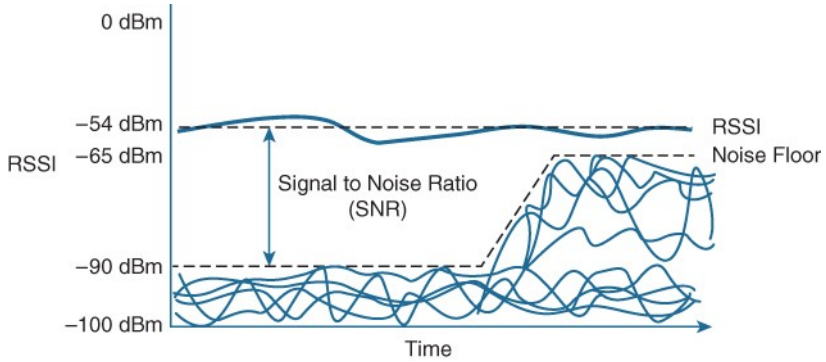

Figure 17-26 *Example of a Changing Noise Floor and SNR*

Carrying Data Over an RF Signal

- Modulation is the process by which a carrier signal is changed in order to carry a data signal.
- Modulation schemes can alter frequency, phase or amplitude of the signal to indicate the zeroes and ones within the data transmission.

Modulation

- To add data to the RF signal, the frequency of the original carrier signal must be preserved. Therefore, there must be some scheme of altering some characteristic of the carrier signal to distinguish a 0 bit from a 1 bit.
- Altering the carrier signal is known as modulation, where the carrier signal is modulated or changed according to some other source. At the receiver, the process is reversed; demodulation interprets the added information based on changes in the carrier signal.
- RF modulation schemes generally have the following goals:
    - Be reasonably immune to interference and noise
    - Be practical to transmit and receive
- Due to the physical properties of an RF signal, a modulation scheme can alter only the following attributes:
    - Frequency, but only by varying slightly above or below the carrier frequency
    - Phase
    - Amplitude

Modulation (Cont.)

- The modulation techniques require some amount of bandwidth centered on the carrier frequency due to the rate of the data being carried and partly due to the overhead from encoding the data and manipulating the carrier signal.
- Narrowband transmissions, such as audio signals over an AM or FM radio, have a relatively low bit rate and little overhead.
- Wireless LANs must carry data at high bit rates, requiring more bandwidth for modulation. Data being sent is spread out across a range of frequencies, known as spread spectrum. Two common spread-spectrum categories:

  Direct sequence spread spectrum (DSSS): Used in the 2.4 GHz band, where a small number of fixed, wide channels support complex phase modulation schemes and somewhat scalable data rates, making it more resilient to disruption.

  Orthogonal Frequency Division Multiplexing (OFDM): Used in both 2.4 and 5 GHz bands, where a single 20 MHz channel contains data that is sent in parallel over multiple frequencies. Each channel is divided into many subcarriers (also called subchannels or tones); both phase and amplitude are modulated with quadrature amplitude modulation (QAM) to move the most data efficiently.

*802.11ax is designed to work on any band from 1 to 7 GHz, provided that the band is approved for use.

Maintaining AP - Client Compatibility

Each step in the 802.11 evolution involves an amendment to the standard, defining things like modulation and coding schemes that are used to carry data over the air. A summary of common amendments to the 802.11 standardis shown in Table17-4.

Table 17-4 A Summary of Common 802.11 Standard Amendments

Standard	2.4 GHz?	5 GHz?	Data Rates Supported	Channel Widths Supported
802.11b	Yes	No	1, 2, 5.5, and 11 Mbps	22 MHz
802.11g	Yes	No	6, 9, 12, 18, 24, 36, 48, and 54 Mbps	22 MHz
802.11a	No	Yes	6, 9, 12, 18, 24, 36, 48, and 54 Mbps	20 MHz
802.11n	Yes	Yes	Up to 150 Mbps* per spatial stream, up to 4 spatial streams	20 or 40 MHz
802.11ac	No	Yes	Up to 866 Mbps per spatial stream, up to 4 spatial streams	20, 40, 80, or 160 MHz
802.11ax	Yes*	Yes*	Up to 1.2 Gbps per spatial stream, up to 8 spatial streams, „High Efficiency WLAN“, band 1 – 7.250 GHz	20, 40, 80, or 160 MHz
802.11ay	60 GHz		20 Gb/s, “Next Generation 60GHz”	Final approval 03/2021
802.11be	Yes*	Yes*	?, “Extremely High Throughput”, band 1 – 7.250 GHz	Final approval 5/2024

*) 802.11ax and 802.11be are designed to work on any band from 1 to 7 GHz.
- Backward compatibility with existing standards
- New bands.

Maintaining AP-Client Compatibility

Newer Wi-Fi Standards include the following:

- The 802.11n amendment was published in 2009 in an effort to scale wireless LAN performance to a theoretical maximum of 600 Mbps. The amendment was unique because it defined a number of additional techniques known as high throughput (HT) that can be applied to either the 2.4 or 5 GHz band.
- The 802.11ac amendment was introduced in 2013 and brought even higher data rates through more advanced modulation and coding schemes, wider channel widths, greater data aggregation during a transmission, and so on. 802.11ac is known as very high throughput (VHT) wireless and can be used only on the 5 GHz band.
- The 802.11ax amendment, also known as Wi-Fi 6 and high efficiency wireless, aims to change the principle that only one device can claim airtime at a time by permitting multiple devices to transmit during the same window of air time. This becomes important in areas that have a high density of wireless devices, all competing for air time and throughput. 802.11ax uses OFDM Access (OFDMA) to schedule and control access to the wireless medium, with channel air time allocated as resource units that can be used by multiple devices simultaneously.

Using Multiple Radios to Scale Performance

Before 802.11n, wireless devices used a single transmitter and a single receiver. In other words, the components formed one radio, resulting in a single radio chain. This is also known as a single-in, single-out (SISO) system.

One secret to the better performance of 802.11n, 802.11ac, and 802.11ax is the use of multiple radio components, forming multiple radio chains. This is known as a multiple-input, multiple-output (MIMO) system.

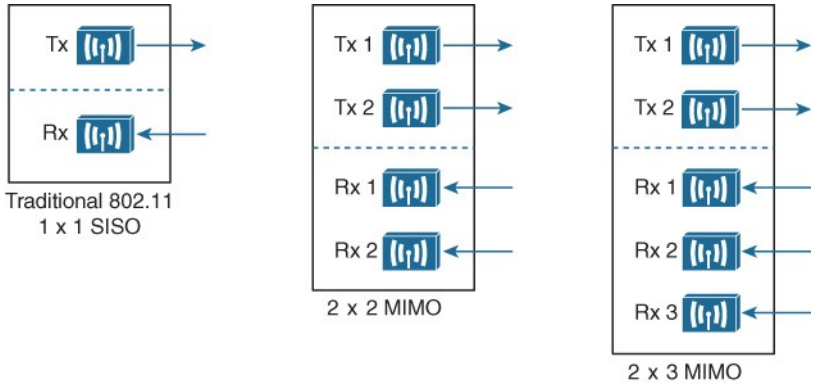

Figure 17-28 *Examples of SISO and MIMO Devices*

Spatial Multiplexing

- To increase data throughput, data can be multiplexed or distributed across two or more radio chains—all operating on the same channel, but separated through spatial diversity. This is known as spatial multiplexing.
- Spatial multiplexing requires a good deal of digital signal processing on both the transmitting and receiving ends. This pays off by increasing the throughput over the channel; the more spatial streams that are available, the more data that can be sent over the channel.
- When the sender and receiver have mismatched spatial stream support, they negotiate the wireless connection and use the lowest number of special streams that they have in common.

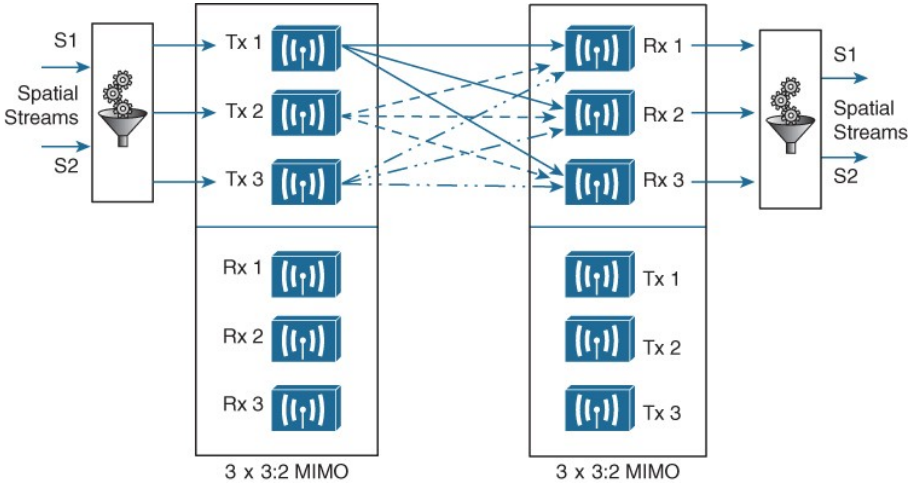

Figure 17-29 *Spatial Multiplexing Between Two 3×3:2 MIMO Devices*

Transmit Beamforming

The 802.11n, 802.11ac, and 802.11ax amendments offer a method to customize the transmitted signal to prefer one receiver over others. By leveraging MIMO, the same signal can be transmitted over multiple antennas to reach specific client locations more efficiently.

With *transmit beamforming (T×BF)*, the phase of the signal is altered as it is fed into each transmitting antenna so that the resulting signals will all arrive in phase at a specific receiver. Figure 17-30 shows a device on the left using transmit beamforming to target device B on the right.

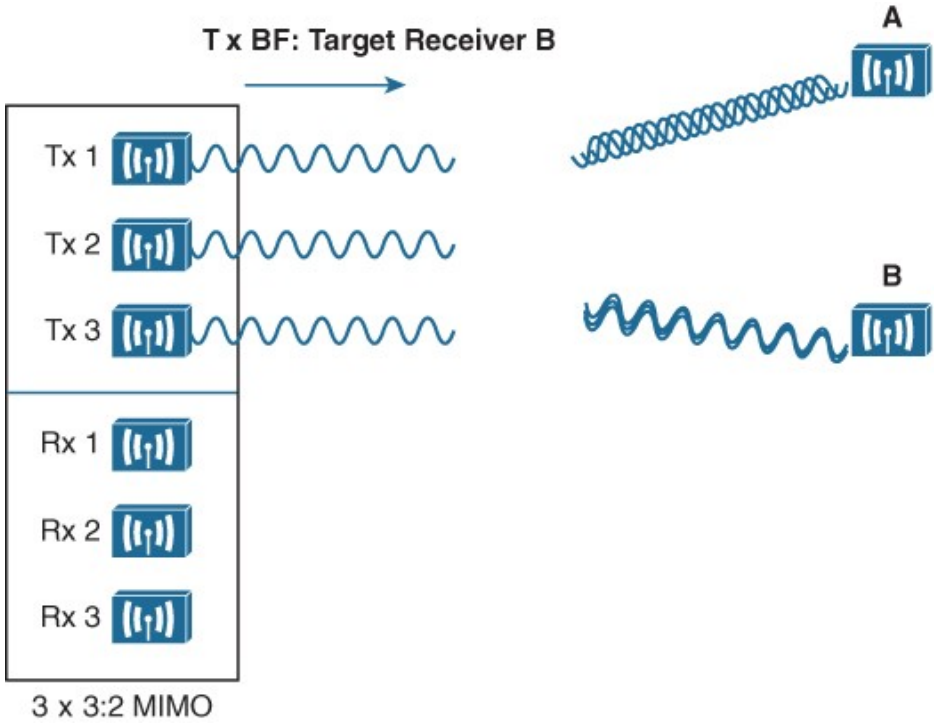

Figure 17-30 *Using Transmit Beamforming to Target a Specific Receiving Device*

Maximal - Ration Combining

Výběrový příjem ☺

When an RF signal is received on a device, it may be degraded or distorted due to a variety of conditions. If that same signal was transmitted over multiple antennas, as in the case of a MIMO device, then the receiving device can attempt to restore it to its original state.

The receiving device can use multiple antennas and radio chains to receive the multiple transmitted copies of the signal. One copy might be better than the others, or one copy might be better for a time, and then become worse than the others. Maximal-ratio combining (MRC) can combine the copies to produce one signal that represents the best version at any given time.

Dynamic Rate Shifting

Figure 17-31 illustrates dynamic rate shifting (DRS) operation on the 2.4 GHz band.

Each concentric circle represents the range supported by a particular modulation and coding scheme. Notice that the white circles denote OFDM modulation (802.11g), and the shaded circles contain DSSS modulation (802.11b).

Each move of the receiver away from the transmitter, into a larger concentric circle, causes a dynamic shift to a reduced data rate, in an effort to maintain the data integrity to the outer reaches of the transmitter's range.

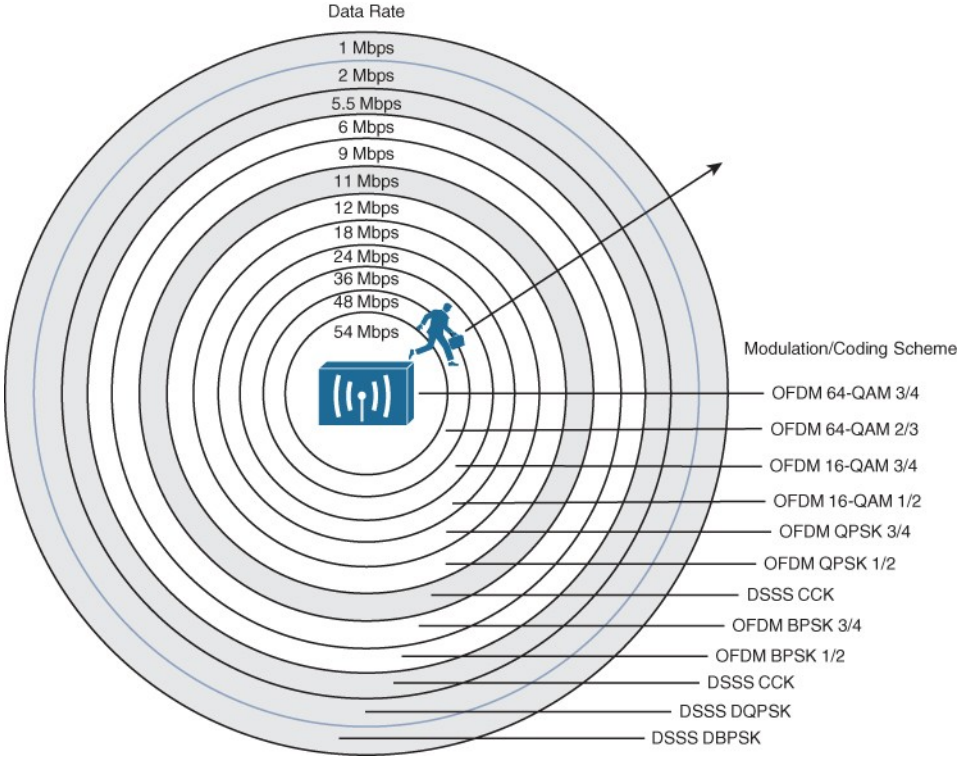

Figure 17-31 *Dynamic Rate Shifting as a Function of Range*

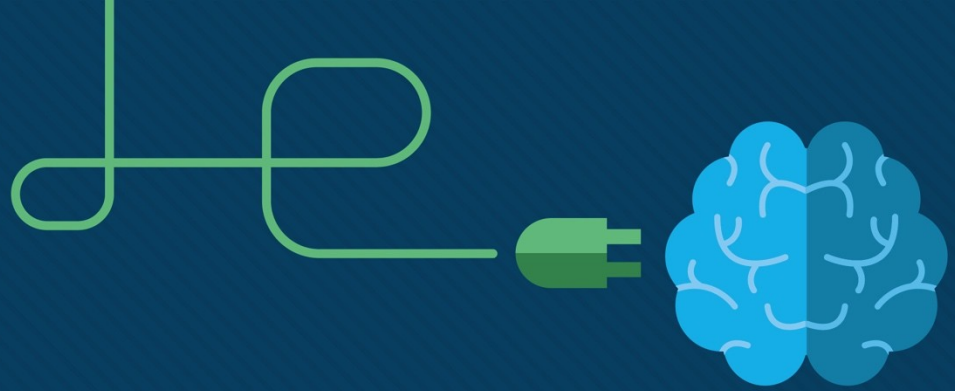

Chapter 18: Wireless Infrastructure

Instructor Materials

CCNP Enterprise: Core Networking

Wireless LAN Topologies

- This chapter looks beyond a single AP to discuss the topologies that can be built with many APs.
- The chapter also discusses the types of antennas you can connect to an AP to provide wireless coverage for various areas and purposes.
- Finally, this chapter discusses how lightweight APs discover and join with wireless LAN controllers in an enterprise network.

AP Modes

Cisco APs can operate in one of two modes:

- **Autonomous** - are self-sufficient and standalone
- **Lightweight** - can support several different network topologies, depending on where the companion wireless LAN controllers (WLCs) are located

Autonomous Topology

Autonomous APs are self-contained, offering one or more standalone basic service sets (BSSs). They are an extension of a switched network, connecting wireless SSIDs to wired VLANs at the access layer.

Fig. 18-1, autonomous APs present two wireless LANs with SSIDs wlan100 and wlan200 to the wireless users. The APs also forward traffic between the wireless LANs and two wired VLANs 100 and 200.

An autonomous AP must also be configured with a management IP address and management VLAN to enable remote management of the AP.

Key Topic

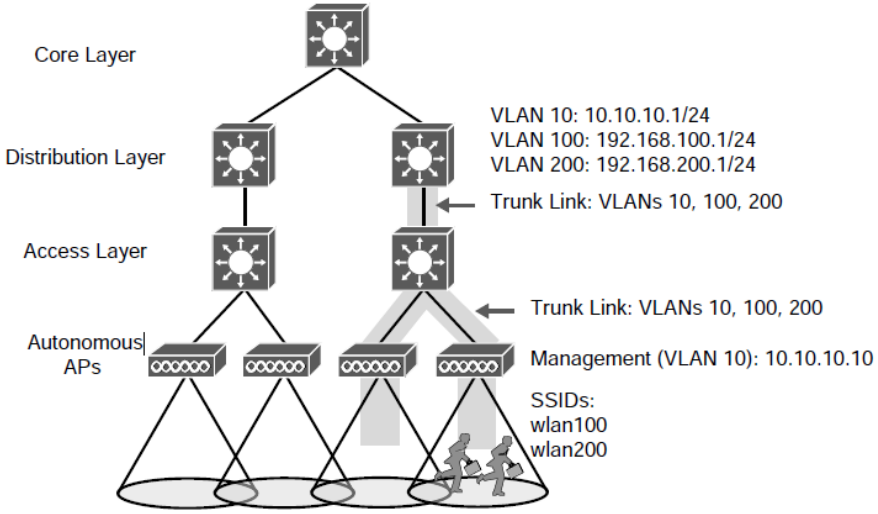

Figure 18-1 *Wireless Network Topology Using Autonomous APs*

Each AP must be configured and maintained individually unless you leverage a management platform such as Cisco Prime Infrastructure.

Autonomous Topology (Cont.)

Because the data and management VLANs may need to reach every autonomous AP, the network configuration and efficiency can become cumbersome as the network scales.

For example, you will likely want to offer the same SSID on many APs so that wireless clients can associate with that SSID in most any location or while roaming between any two APs.

You may want to extend the VLAN and IP subnet to each and every AP so that clients do not have to request a new IP address for each new association.

A topology using autonomous APs does have one nice feature: a short and simple path for data to travel between the wireless and wired networks.

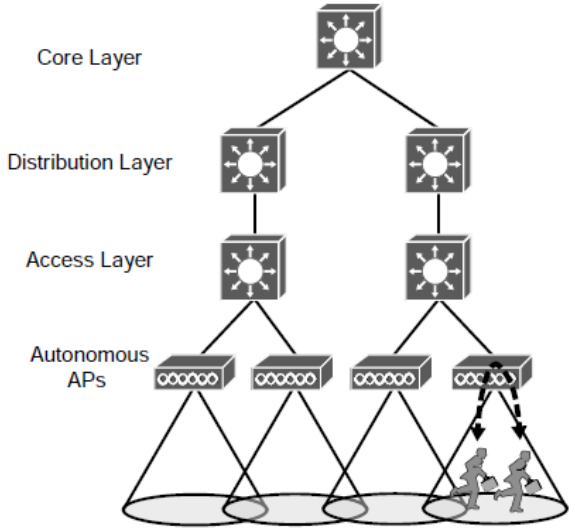

Figure 18-2 *Shortest Data Path Through an Autonomous AP Topology*

In Figure 18-2, two wireless users are associated to the same autonomous AP. One can reach the other through the AP, without having to pass up into the wired network. This is not always the case with lightweight AP topologies.

Lightweight AP Topologies

CAPWAP = Control And Provisioning of Wireless Access Points

In lightweight mode, an AP **loses** its self-sufficiency to provide a working BSS for wireless users. It has to join a WLC to become fully functional.

This is known as a split-MAC architecture, where the AP handles most of the realtime 802.11 processes and the WLC performs the management functions.

An AP and a WLC are joined by a logical pair of CAPWAP tunnels that extend through the wired network infrastructure. Control and data traffic are transported across the tunnels.

Several topologies can be built from a WLC and a collection of APs. These differ according to where the WLC is located within the network.

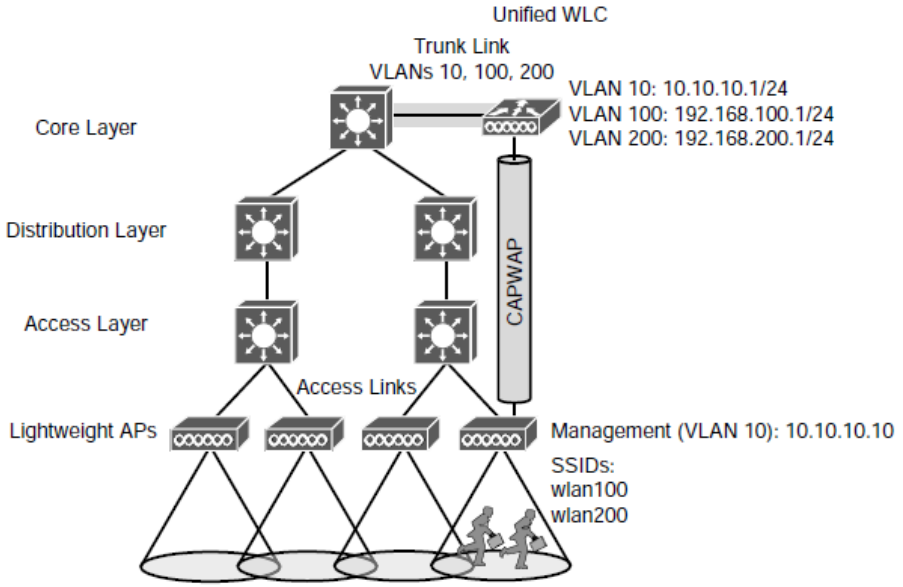

Figure 18-3 *WLC Location in a Centralized Wireless Network Topology*

Fig. 18-3, a WLC is placed in a central location, so it can maximize the number of APs joined to it. This is known as a centralized or unified wireless LAN topology. Each AP has its own CAPWAP tunnel to the WLC.

Lightweight AP Topologies - Centralized

A Cisco unified WLC meant for a large enterprise can support up to 6000 APs.

The Layer 3 boundary for each data VLAN is handled at or near the WLC, so the VLANs need only exist at that location, indicated by the shaded link.

Each AP still has its own unique management IP address, but it connects to an access layer switch via an access link rather than a trunk link. Even if multiple VLANs and WLANs are involved, they are carried over the same CAPWAP tunnel to and from the AP. Therefore, the AP needs only a single IP address to terminate the tunnel.

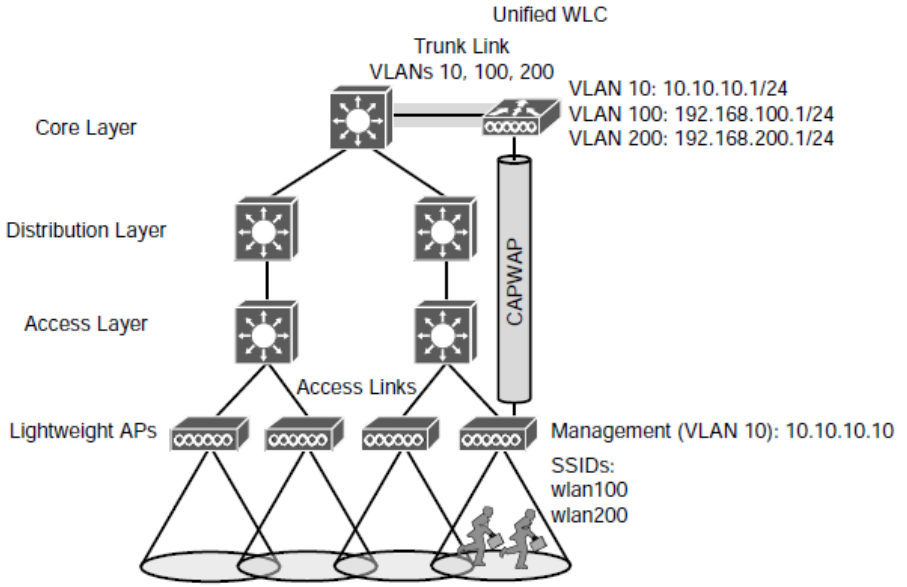

Figure 18-3 *WLC Location in a Centralized Wireless Network Topology*

As a wireless user moves through the coverage areas of the four APs, he might associate with many different APs in the access layer. Because all of the APs are joined to a single WLC, that WLC can easily maintain the user's connectivity to all other areas of the network as he moves around.

Lightweight AP Topologies – Centralized (Cont.)

The traffic from one client must pass through the AP, where it is encapsulated in the CAPWAP tunnel, and then travel high up into the network to reach the WLC, where it is unencapsulated and examined. The process then reverses.

The length of the tunnel path can be a great concern for lightweight APs.

The round-trip time (RTT) between an AP and a controller should be less than 100 ms so that wireless communication can be maintained in near real time. If the path has more latency than that, the APs may decide that the controller is not responding fast enough, so they may disconnect and find another, more responsive controller.

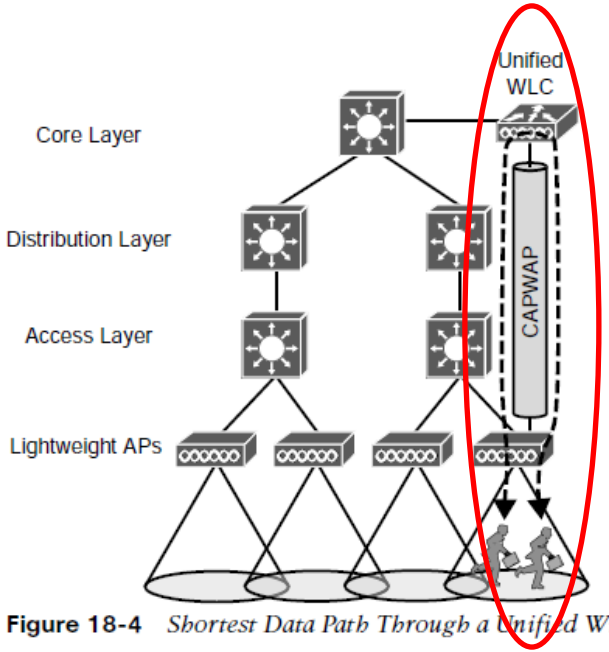

Figure 18-4 *Shortest Data Path Through a Unified Wireless Network Topology*

Lightweight AP Topologies – Embedded Wireless Topology

A WLC can be located further down in the network hierarchy.

Fig. 18-5, the WLC is co-located with an access layer switch. This is known as an embedded wireless network topology because the WLC is embedded in the switch hardware.

With user access merged into one layer, it becomes easier to apply common access and security policies. Notice that each AP connects to an access switch for network connectivity as well as split-MAC functionality, so the CAPWAP tunnel becomes really short.

The embedded topology can be cost-effective because the same switching platform is used for both wired and wireless purposes. Ideally, each access layer switch would have its own embedded WLC. A Cisco embedded WLC typically supports up to 200 APs.

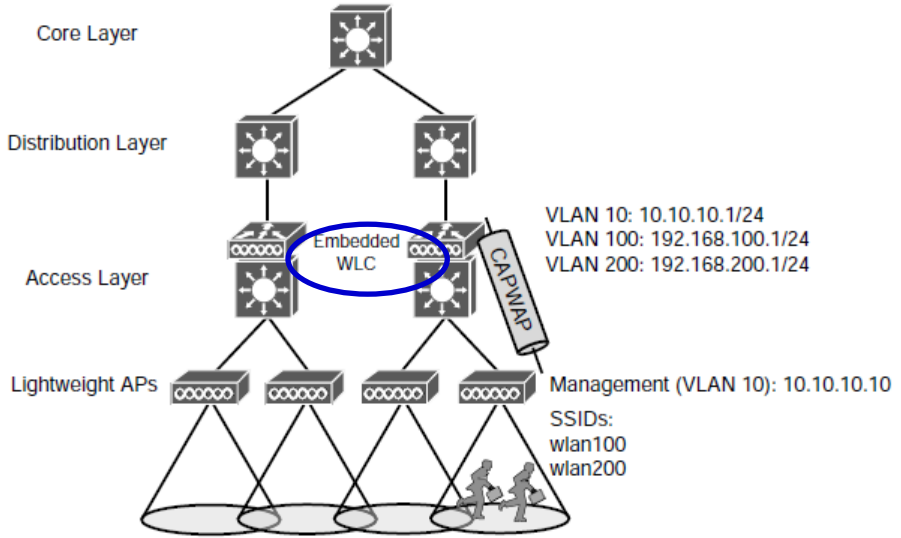

Figure 18-5 *WLC Location in an Embedded Wireless Network Topology*

Lightweight AP Topologies – Embedded Wireless Topology (Cont.)

If the CAPWAP tunnel is relatively short in an embedded topology, that must mean wireless devices can reach each other more efficiently.

Fig. 18-6, shows, the traffic path from one user to another must pass through an AP, the access switch (and WLC), and back down through the AP.

In contrast, traffic from a wireless user to a central resource such as a data center or the internet travels through the CAPWAP tunnel, is unencapsulated at the access layer switch (and WLC), and travels normally up through the rest of the network layers.

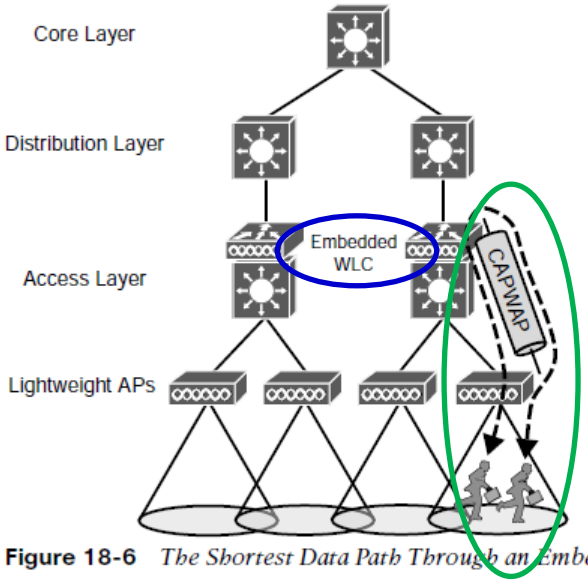

Figure 18-6 *The Shortest Data Path Through an Embedded Wireless Network Topology*

Lightweight AP Topologies – Mobility Express Network Topology

It is also possible to move the WLC even below the access layer and into an AP.

Fig. 18-7, illustrates the Mobility Express topology, where a fully functional Cisco AP also runs software that acts as a WLC. This can be useful in small scale environments, such as small, midsize, or multi-site branch locations, where you might not want to invest in dedicated WLCs at all.

The AP that hosts the WLC forms a CAPWAP tunnel with the WLC, as do any other APs at the same location. A Mobility Express WLC can support up to 100 APs.

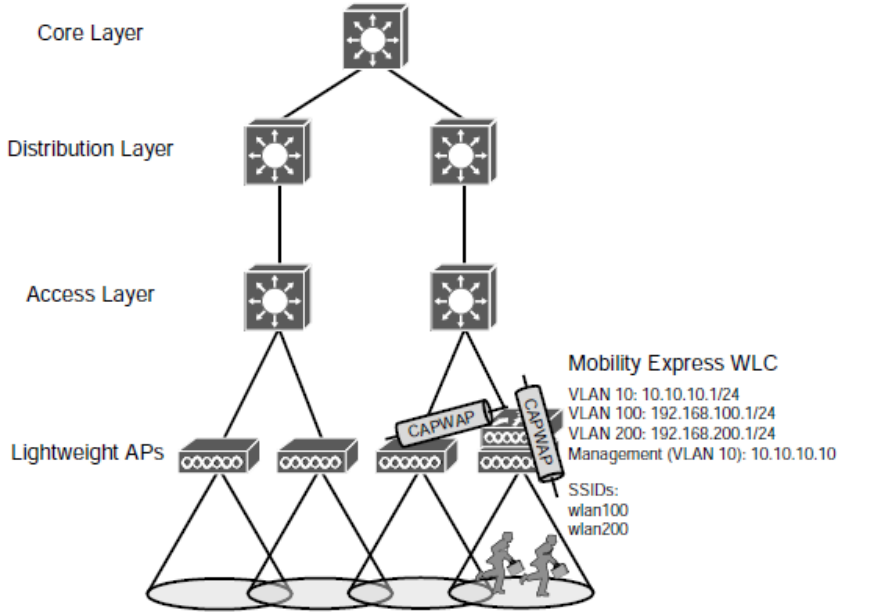

Figure 18-7 *WLC Location in a Mobility Express Wireless Network Topology*

Pairing Lightweight APs and WLCs

- A Cisco lightweight wireless AP needs to be paired with a WLC to function.
- Each AP must discover and bind itself with a controller before wireless clients can be supported.
- Cisco lightweight APs are designed to be “touch free,” but you have to configure the switch port, where the AP connects, with the correct access VLAN, access mode, and inline power settings, then the AP can power up and use a variety of methods to find a viable WLC to join.

AP States

A lightweight AP goes through a variety of states defined as part of the Control and Provisioning of Wireless Access Points (CAPWAP) specification. The AP enters the states in a specific order; the sequence of states is called a state machine:

1. **AP boots -** Once an AP receives power, it boots on a small IOS image so that it can work through the remaining states and communicate over its network connection. The AP must also receive an IP address from either a DHCP server or a static configuration so that it can communicate over the network.

2. **WLC discovery -** The AP goes through a series of steps to find one or more controllers that it might join.

3. **CAPWAP tunnel -** The AP attempts to build a CAPWAP tunnel with one or more controllers. The tunnel will provide a secure Datagram Transport Layer Security (DTLS) channel for subsequent AP-WLC control messages. The AP and WLC authenticate each other through an exchange of digital certificates.

4. **WLC join -** The AP selects a WLC from a list of candidates and then sends a CAPWAP Join Request message to it. The WLC replies with a CAPWAP Join Response message.

5. **Download image -** The WLC informs the AP of its software release. If the AP's own software is a different release, the AP downloads a matching image from the controller, reboots to apply the new image, and then returns to step 1.

AP States (Cont.)

6. **Download config -** The AP pulls configuration parameters down from the WLC and can update existing values with those sent from the controller. Settings include RF, service set identifier (SSID), security, and quality of service (QoS) parameters.

7. **Run state -** Once the AP is fully initialized, the WLC places it in the "run" state. The AP and WLC then begin providing a BSS and begin accepting wireless clients.

8. **Reset -** If an AP is reset by the WLC, it tears down existing client associations and any CAPWAP tunnels to WLCs. The AP then reboots and starts through the entire state machine again.

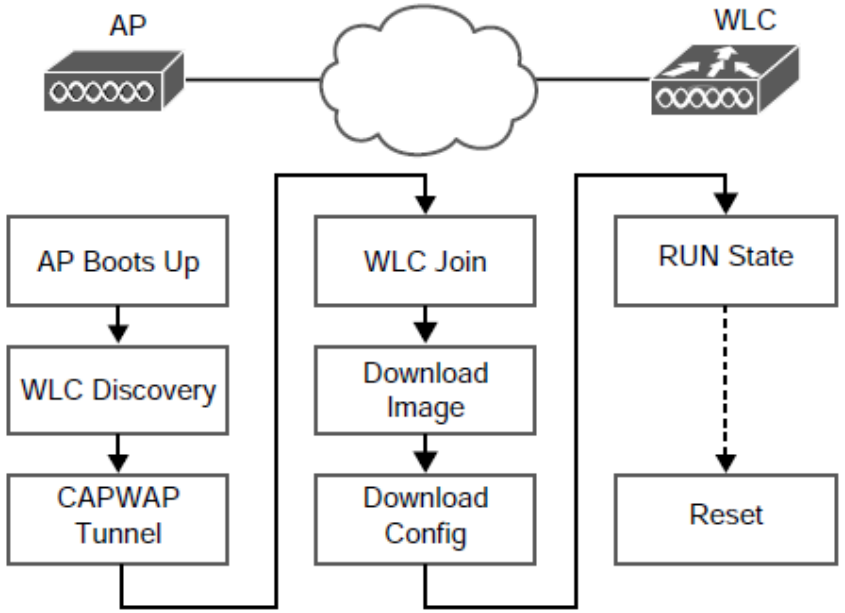

Figure 18-8 *State Machine of a Lightweight AP*

If there is a chance an AP could rehome with another WLC, you should make sure that both WLCs are running the same code release. Otherwise, the AP move should happen at a planned time, like during a maintenance window. You can predownload a new release to the controller's APs prior to rebooting the WLC.

Discovering a WLC

To discover a WLC, an AP sends a unicast CAPWAP Discovery Request to a controller's IP over UDP port 5246 or a broadcast to the local subnet. If the controller exists, it returns a CAPWAP Discovery Response to the AP.

An AP must discover any WLCs that it can join without any preconfiguration. Several methods of discovery are used and the sequence of discovery is as follows:

1. The AP broadcasts a CAPWAP Discovery Request on its local wired subnet. Any WLCs on the subnet answer with a CAPWAP Discovery Response.

2. An AP can be "primed" with up to 3 controllers: a primary, a secondary, and a tertiary. These are stored in NVRAM so that the AP can remember them after a reboot. Otherwise, if an AP has previously joined a WLC, it may have stored up to 8 out of a list of 32 WLC addresses that it received from the last controller it joined. The AP attempts to contact as many controllers as possible to build a list of candidates.

3. The DHCP server that supplies an IP can also send DHCP option 43 to suggest WLC addresses.

4. The AP attempts to resolve the name CISCO-CAPWAP-CONTROLLER.localdomain with a DNS request (where localdomain is the domain name learned from DHCP). If the name resolves to an IP address, the controller attempts to contact a WLC at that address.

5. If none of the steps has been successful, the AP resets itself and restarts the discovery process again.

Discovering a WLC (Cont.)

If the AP and controllers lie on different subnets, you can configure the local router to relay any broadcast requests on UDP port 5246 to specific controller addresses.

Use the following configuration commands:

router(config)# **ip forward-protocol udp 5246**

router(config)# **interface vlan *number***

router(config-int)# **ip helper-address *WLC1-MGMT-ADDR***

router(config-int)# **ip helper-address *WLC2-MGMT-ADDR***

Selecting a WLC

Joining a WLC involves sending it a CAPWAP Join Request and waiting for it to return a CAPWAP Join Response. From that point on, the AP and WLC build a DTLS tunnel to secure their CAPWAP control messages.

The WLC selection process consists of the following three steps:

1. If the AP has previously joined a controller and has been configured or “primed” with a primary, secondary, and tertiary controller, it tries to join those controllers in succession.

2. If the AP does not know of any candidate controller, it tries to discover one. If a controller has been configured as a master controller, it responds to the AP’s request.

3. The AP attempts to join the least-loaded WLC, to load balance APs across a set of controllers. During the discovery phase, each controller reports its load—the ratio of the number of currently joined APs to the total AP capacity.

The least-loaded WLC is the one with the lowest ratio. If the controller already has the maximum number of APs joined to it, it rejects any additional APs.

To provide flexibility in supporting APs on an oversubscribed controller, you can configure the APs with a priority value. Once a controller is full of APs, it rejects an AP with the lowest priority to make room for a new one that has a higher priority.

Maintaining WLC Availability

If a controller full of 1000 APs fails, all 1000 APs must detect the failure, discover other controllers, and then select the least-loaded one to join. During that time, wireless clients can be left stranded with no connectivity.

The most deterministic approach is to use the primary, secondary, and tertiary controller fields in every AP.

Once an AP joins a controller, it sends keepalive messages to the controller over the wired network. By default, keepalives are sent every 30 seconds. If a keepalive is not answered, an AP escalates by sending four more keepalives at 3-second intervals. If it does not answer, the AP presumes that the controller has failed. The AP then moves quickly to find a successor to join.

Using default values, an AP can detect controller failure in 35 seconds. Using minimum values, failure can be detected in only 6 seconds.

WLCs also support high availability (HA) with stateful switchover (SSO) redundancy. One controller takes on the active role and the other a hot standby mode. The APs only need to know the active primary controller.

The active unit keeps CAPWAP tunnels, AP states, client states, configurations, and image files all in sync with the hot standby unit. The active controller also synchronizes the state of each associated client that is in the RUN state with the hot standby controller. If the active controller fails, the standby will already have the current state information for each AP and client, making the failover process transparent to the end users.

Cisco AP Modes

From the WLC, you can configure a lightweight AP to operate in one of the following modes:

- **Local** - The default lightweight mode that offers one or more functioning BSSs on a specific channel. During times when it is not transmitting, the AP scans the other channels to measure the level of noise, measure interference, discover rogue devices, and match against intrusion detection system (IDS) events.
- **Monitor** - The AP does not transmit at all, but its receiver is enabled to act as a dedicated sensor. The AP checks for IDS events, detects rogue access points, and determines the position of stations through location-based services.
- **FlexConnect** - An AP at a remote site can locally switch traffic between an SSID and a VLAN if its CAPWAP tunnel to the WLC is down and if it is configured to do so.
- **Sniffer** - An AP dedicates its radios to receiving 802.11 traffic from other sources, much like a sniffer or packet capture device. The captured traffic is then forwarded to a PC running network analyzer software such as LiveAction Omnipeek or Wireshark, where it can be analyzed further.
- **Rogue detector** - An AP dedicates itself to detecting rogue devices by correlating MAC addresses heard on the wired network with those heard over the air. Rogue devices are those that appear on both networks.

Cisco AP Modes (Cont.)

- **Bridge** - An AP becomes a dedicated bridge (point-to-point or point-to-multipoint) between two networks. Two APs in bridge mode can be used to link two locations separated by a distance. Multiple APs in bridge mode can form an indoor or outdoor mesh network.
- **Flex+Bridge** - FlexConnect operation is enabled on a mesh AP.
- **SE-Connect** - The AP dedicates its radios to spectrum analysis on all wireless channels. You can remotely connect a PC running software such as MetaGeek Chanalyzer or Cisco Spectrum Expert to the AP to collect and analyze the spectrum analysis data to discover sources of interference.

A lightweight AP is normally in local mode when it is providing BSSs and allowing client devices to associate to wireless LANs. When an AP is configured to operate in one of the other modes, local mode (and the BSSs) is disabled.

Leveraging Antennas for Wireless Coverage

- One type of antenna cannot fit every application.
- Antennas come in many sizes and shapes, each with its own gain value and intended purpose.
- The following section describes antenna characteristics in more detail.

Radiation Patterns

- Antenna gain is normally a comparison of one antenna against an isotropic antenna and is measured in dBi (decibel-isotropic).
- An isotropic antenna does not actually exist because it is ideal, perfect, and impossible to construct.
- An isotropic antenna is shaped like a tiny round point.
- When an alternating current is applied, an RF signal is produced, and the electromagnetic waves are radiated equally in all directions.
- The energy produced by the antenna takes the form of an ever-expanding sphere.
- A plot that shows the relative signal strength around an antenna is known as the radiation pattern.

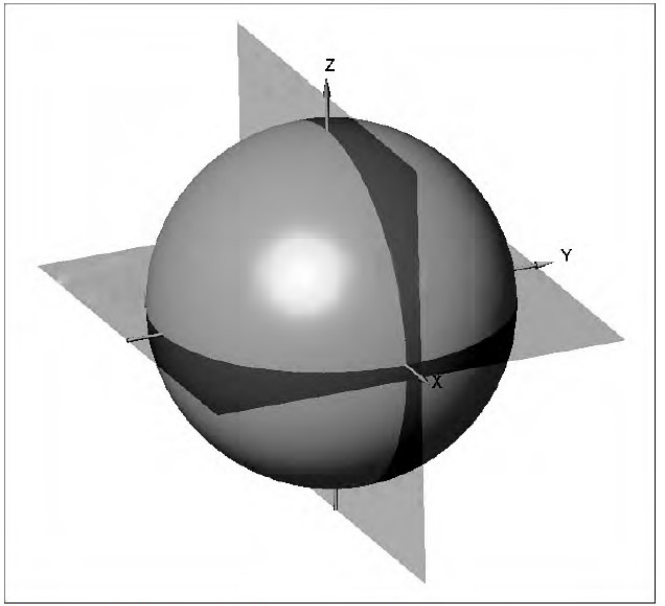

Figure 18-9 *Plotting the Radiation Pattern of an Isotropic Antenna*

Radiation Patterns (Cont.)

- The XY plane, which lies flat along the horizon, is known as the H plane, or the horizontal (azimuth) plane.
- The XZ plane, which lies vertically along the elevation of the sphere, is known as the E plane, or elevation plane.
- The outline of each plot can be recorded on a polar plot.
- The outermost circle usually represents the strongest signal strength, and the inner circles represent weaker signal strength.
- The antenna is placed at the center of the polar plots.
- As you decide to place APs in their actual locations, you might have to look at various antenna patterns and try to figure out whether the antenna is a good match for the environment you are trying to cover with an RF signal.

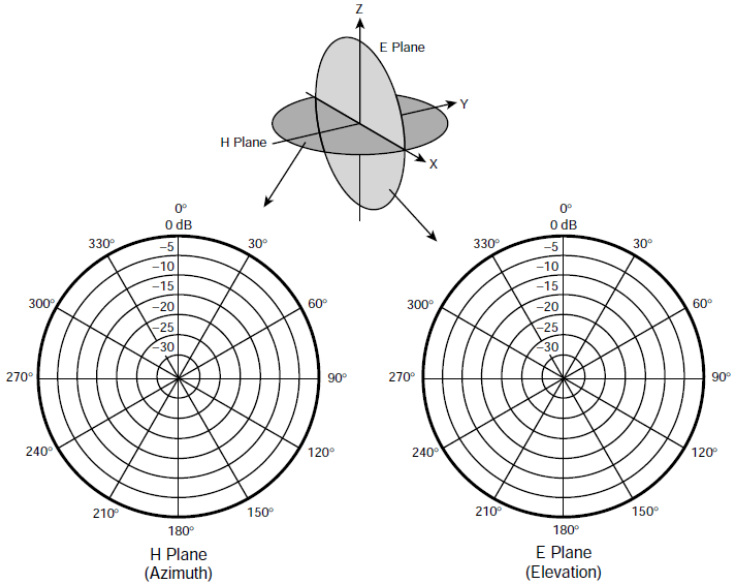

Figure 18-10 *Recording an Isotropic Antenna Pattern on E and H Polar Plots*

Gain

- Antenna amplify or add gain to the signal by shaping the RF energy as it is propagated into free space. The gain of an antenna is a measure of how effectively it can focus RF energy in a certain direction.
- Think of a zero gain antenna producing a perfect sphere. If the sphere is made of rubber, you could press on it in various locations and change its shape. As the sphere is deformed, it expands in other directions. Figure 18-11 shows some simple examples, along with some examples of gain values.
- The gain is lower for omnidirectional antennas, which are made to cover a widespread area, and higher for directional antennas, which are built to cover more focused areas.
- The gain is typically not indicated on either E or H plane radiation pattern plots. The only way to find an antenna's gain is to look at the manufacturer's specifications.

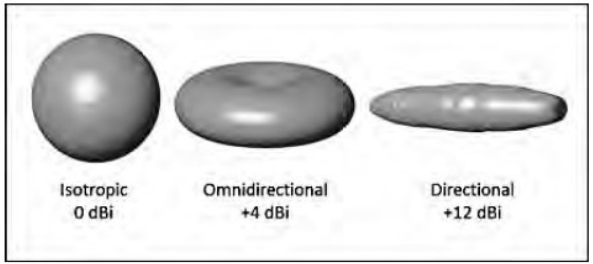

Figure 18-11 *Radiation Patterns for the Three Basic Antenna Types*

Beamwidth

Many manufacturers list the beamwidth of an antenna as a measure of the antenna's focus.

Beamwidth is normally listed in degrees for both the H and E planes.

The beamwidth is determined by finding the strongest point on the plot, which is usually somewhere on the outer circle. Next, the plot is followed in either direction until the value decreases by 3 dB, indicating the point where the signal is one-half the strongest power.

A line is drawn from the center of the plot to intersect each 3 dB point, and then the angle between the two lines is measured.

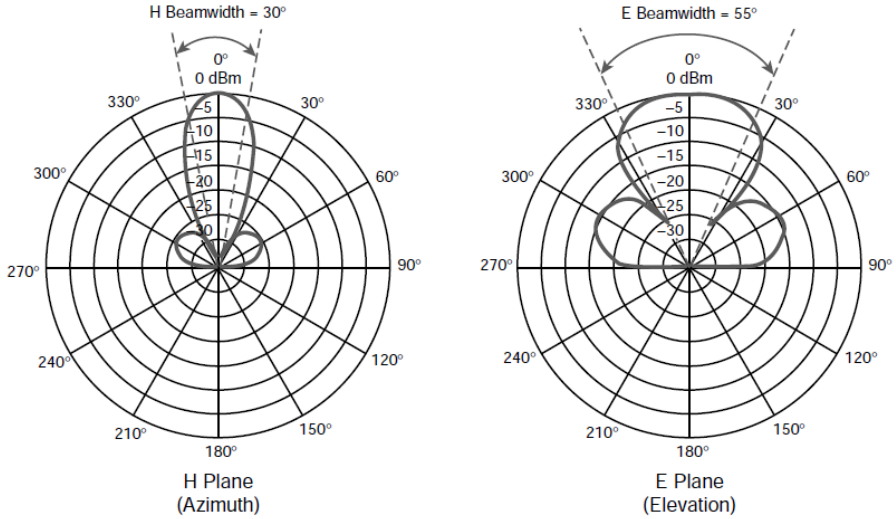

Figure 18-12 *Example of Antenna Beamwidth Measurement*

Figure 18-12 shows a simple example. The H plane has a beamwidth of 30 degrees, and the E plane has a beamwidth of 55 degrees.

Polarization

A wave has two components: an electrical field wave and a magnetic field wave.

The electrical portion of the wave will always leave the antenna in a certain orientation. If the wire is pointing vertically it will produce a wave that oscillates up and down in a vertical direction.

The electrical field wave's orientation is called the antenna polarization.

Antennas that produce vertical oscillation are vertically polarized; those that produce horizontal oscillation are horizontally polarized.

Antenna polarization at the transmitter must be matched to the polarization at the receiver. If the polarization is mismatched, the received signal can be severely degraded.

In Fig. 18-13 The transmitter and receiver along the top both use vertical polarization, so the received signal is optimized.
The pair along the bottom is mismatched, causing the signal to be poorly received.

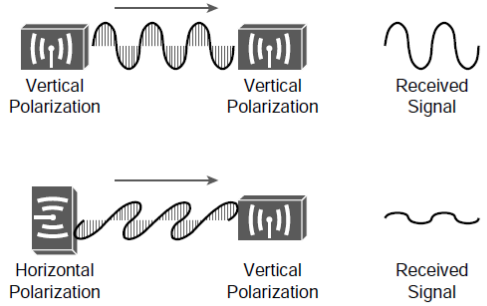

Figure 18-13 *Matching the Antenna Polarization Between Transmitter and Receiver*

Omnidirectional Antennas

An omnidirectional antenna tends to propagate a signal equally in all directions away from the cylinder but not along the cylinder's length.

The result is a donut-shaped pattern that extends further in the H plane than in the E plane.

This type of antenna is well suited for broad coverage of a large room or floor area, with the antenna located in the center.

Because an omnidirectional antenna distributes the RF energy throughout a broad area, it has a relatively low gain.

A common type of omnidirectional antenna is the dipole.

As its name implies, the dipole has two separate wires that radiate an RF signal when an alternating current is applied across them.

Dipoles usually have a gain of around +2 to +5 dBi.

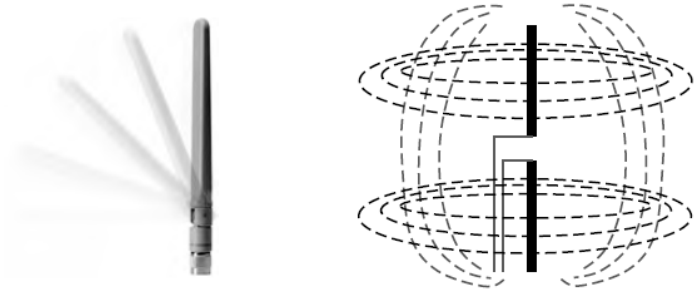

Figure 18-14 *Cisco Dipole Antenna*

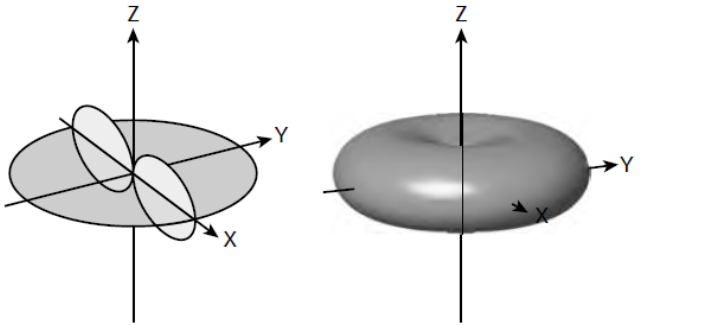

Figure 18-16 *Dipole Radiation Pattern in Three Dimensions*

Omnidirectional Antennas (Cont.)

To reduce the size of an omnidirectional antenna, many Cisco wireless access points (APs) have integrated antennas that are hidden inside the device's smooth case. For example, the AP shown in Figure 18-17 has six tiny antennas hidden inside it.

Figure 18-17 *Cisco Wireless Access Point with Integrated Omnidirectional Antennas*

Directional Antennas

Directional antennas have a higher gain than omnidirectional antennas because they focus the RF energy in one general direction.

Typical applications include elongated indoor areas, such as the rooms along a long hallway or the aisles in a warehouse. They can also be used to cover outdoor areas out away from a building or long distances between buildings.

If they are mounted against a ceiling, pointing downward, they can cover a small floor area to reduce an AP's cell size.

Patch antennas have a flat rectangular shape, as shown in Figure 18-19, so that they can be mounted on a wall or ceiling.

Patch antennas have a typical gain of about 6 to 8 dBi in the 2.4 GHz band and 7 to 10 dBi at 5 GHz.

Figure 18-19 *Typical Cisco Patch Antenna*

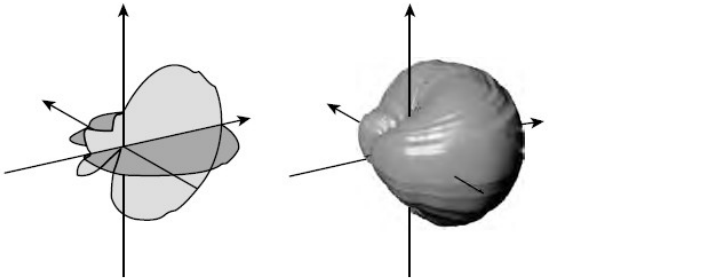

Figure 18-21 *Patch Antenna Radiation Pattern in Three Dimensions*

Yagi Antennas

Figure 18-22 shows the Yagi–Uda antenna, named after its inventors, and more commonly known as the Yagi.

Although its outer case is shaped like a thick cylinder, the antenna is actually made up of several parallel elements of increasing length.

A Yagi produces a more focused egg-shaped pattern that extends out along the antenna's length, as shown in Figure 18-24.

Yagi antennas have a gain of about 10 to 14 dBi.

Figure 18-22 *Cisco Yagi Antenna*

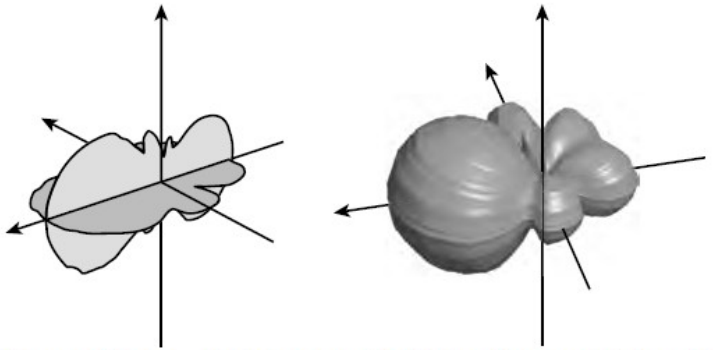

Figure 18-24 *Yagi Antenna Radiation Pattern in Three Dimensions*

Parabolic Dish Antennas

In a line-of-sight wireless path, an RF signal must be propagated a long distance using a narrow beam.

Highly directional antennas focus the RF energy along one narrow elliptical pattern.

Dish antennas, as shown in Fig. 18-25, use a parabolic dish to focus received signals onto an antenna mounted at the center.

The parabolic shape causes any waves arriving from the line of sight will be reflected onto the center antenna element that faces the dish.

Transmitted waves are just the reverse. They are aimed at the dish and reflected such that they are propagated away from the dish along the line of sight.

The focused pattern gives the antenna a gain of between 20 and 30 dBi—the highest of all the wireless antennas.

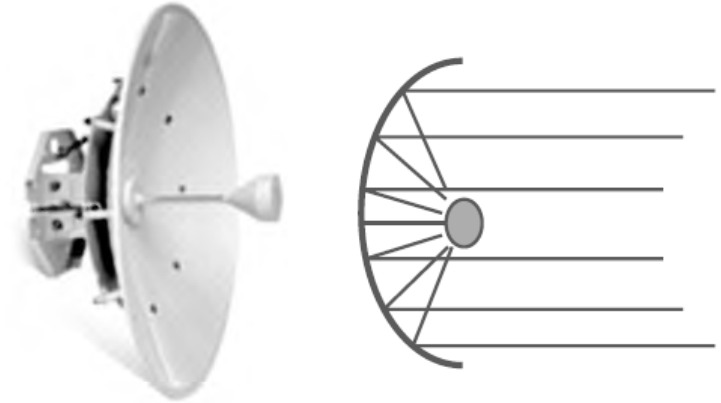

Figure 18-25 *Cisco Parabolic Dish Antenna*

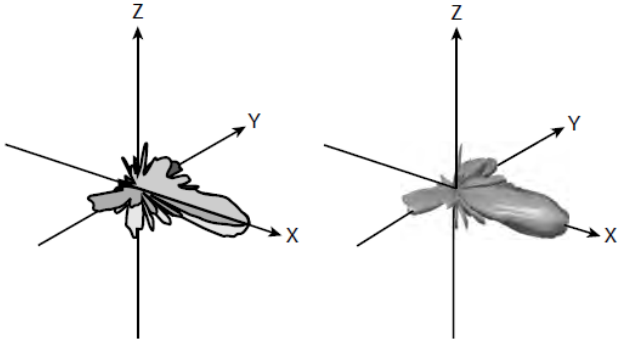

Figure 18-27 *Parabolic Dish Antenna Radiation Pattern in Three Dimensions*

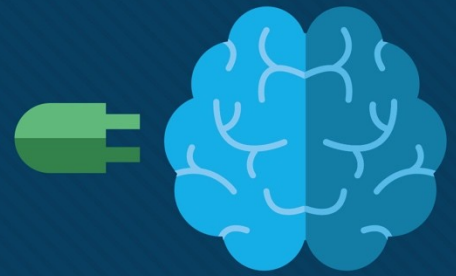

Chapter 19: Understanding Wireless Roam and Location Services

Instructor Materials

CCNP Enterprise: Core Networking

Roaming Overview

- To understand how wireless roaming works, start with simple scenarios such as roaming between access points when no controller is present and when only one controller is present.

Before Roaming Between Autonomous APs

A wireless client must associate and authenticate with an AP before it can use the AP's BSS to access the network. A client can also move from one BSS to another by roaming between APs. The client actively scans channels and sends probe requests to discover candidate APs, and then the client selects one and tries to reassociate with it.

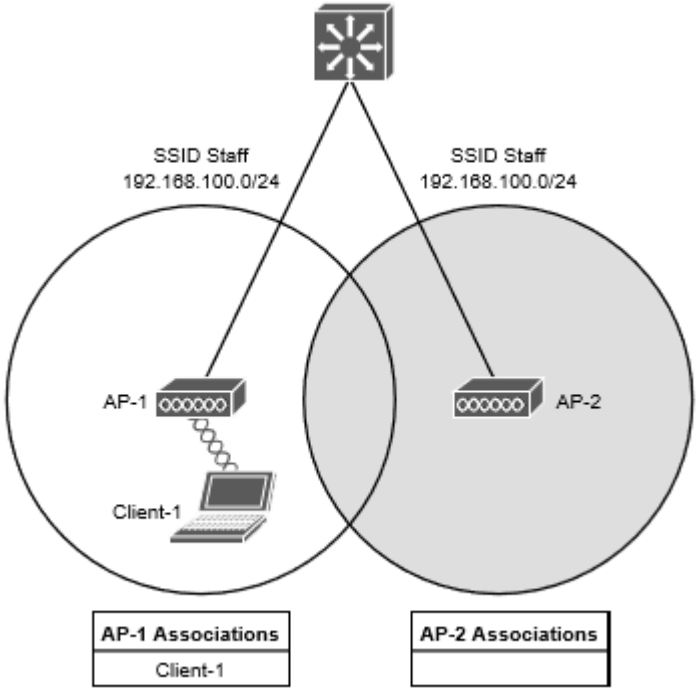

Figure 19-1 *Before Roaming Between Autonomous APs*

After Roaming Between Autonomous APs

The client begins to move into AP 2's cell. Somewhere near the cell boundary, the client decides that the signal from AP 1 has degraded and it should look elsewhere for a stronger signal. The client decides to roam and reassociate with AP 2.

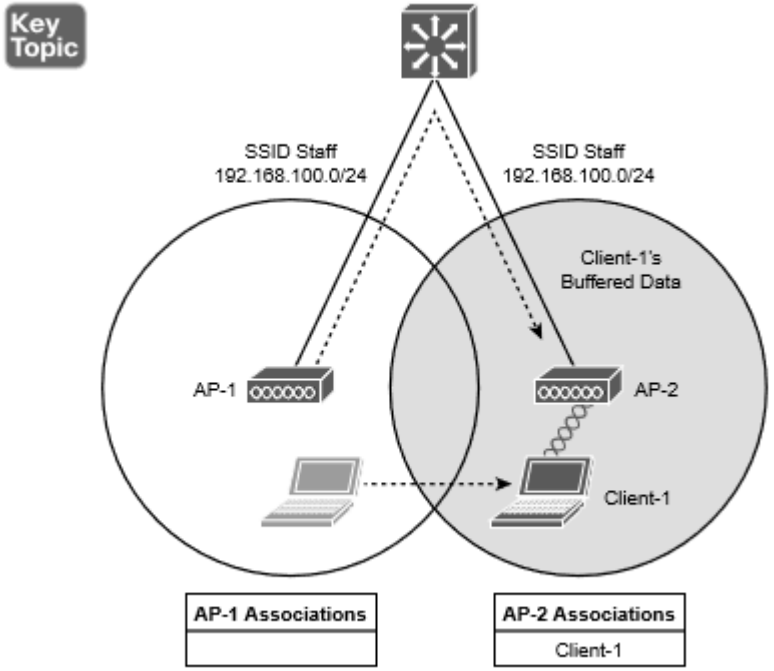

Figure 19-2 *After Roaming Between Autonomous APs*

Successive Roams of a Mobile Client

When a wireless client begins to move, it might move along an arbitrary path. Each time the client decides that the signal from one AP has degraded enough, it attempts to roam to a new, better signal belonging to a different AP and cell. **The exact location of each roam depends on the client's roaming algorithm.** To illustrate typical roaming activity, each roam in Figure 19-3 is marked with a dark ring.

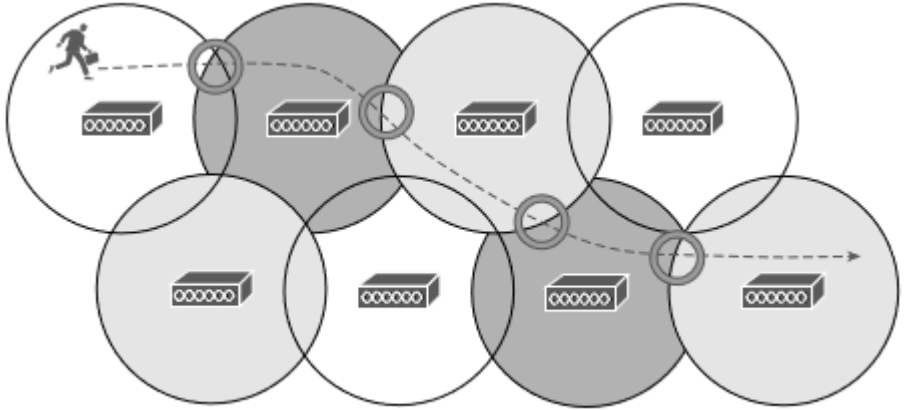

Figure 19-3 *Successive Roams of a Mobile Client*

Intracontroller Roaming

In a Cisco wireless network, lightweight APs are bound to a wireless LAN controller through CAPWAP tunnels. **The controller handles the roaming process, rather than the APs**, because of the split-MAC architecture.

If both APs involved in a client roam are bound to the same controller, the controller has to update its client association table so that it knows which CAPWAP tunnel to use to reach the client.

Key Topic

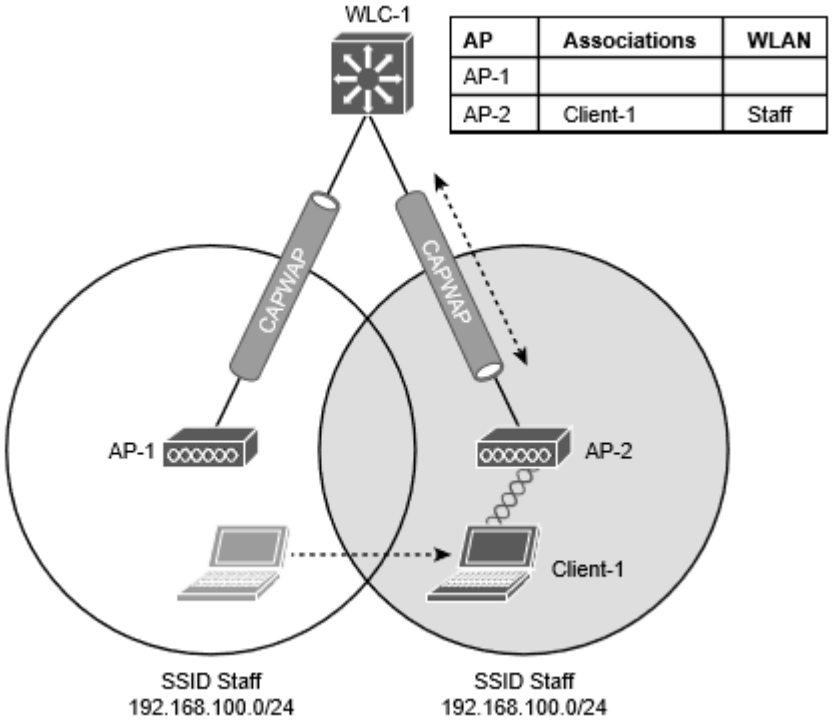

Figure 19-5 *Cisco Wireless Network After an Intracontroller Roam*

Intracontroller Roaming (Cont.)

Efficient roaming is especially important when time-critical applications are being used over the wireless network.

When a roam occurs, there could be a brief time when the client is not fully associated with either AP. So long as that time is held to a minimum, the end user probably will not even notice that the roam occurred.

Along with the client reassociation, a couple other processes can occur:

- **DHCP -** The client may be programmed to renew the DHCP lease on its IP address or to request a new address.
- **Client authentication -** The controller might be configured to use an 802.1x method to authenticate each client on a WLAN.

Cryptographic Key Exchange Techniques

The client authentication process presents a challenge because the dialog between a controller and a RADIUS server, in addition to the cryptographic keys that need to be generated and exchanged, can take a considerable time to accomplish.

Cisco controllers offer three techniques to help streamline this process:

- **Cisco Centralized Key Management (CCKM) -** One controller maintains a database of clients and keys on behalf of its APs and provides them to other controllers and their APs as needed during client roams. CCKM requires Cisco Compatible Extensions (CCX) support from clients.
- **Key caching -** Each client maintains a list of keys used with prior AP associations and presents them as it roams. The destination AP must be present in this list, which is limited to eight AP/key entries.
- **802.11r -** This 802.11 amendment addresses fast roaming or fast BSS transition; a client can cache a portion of the authentication server's key and present that to future APs as it roams. The client can also maintain its QoS parameters as it roams.

Roaming Between Centralized Controllers

- When two or more controllers support the APs in an enterprise, the APs can be distributed across them. As always, when clients become mobile, they roam from one AP to another—except they could also be roaming from one controller to another, depending on how neighboring APs are assigned to the controllers. As a network grows, AP roaming can scale too by organizing controllers into mobility groups.

Before an Intercontroller Roam

When a client roams from one AP to another and those APs lie on two different controllers, the client makes an intercontroller roam.

Figure 19-6 shows a simple scenario prior to a roam. Controller WLC 1 has one association in its database—that of Client 1 on AP 1.

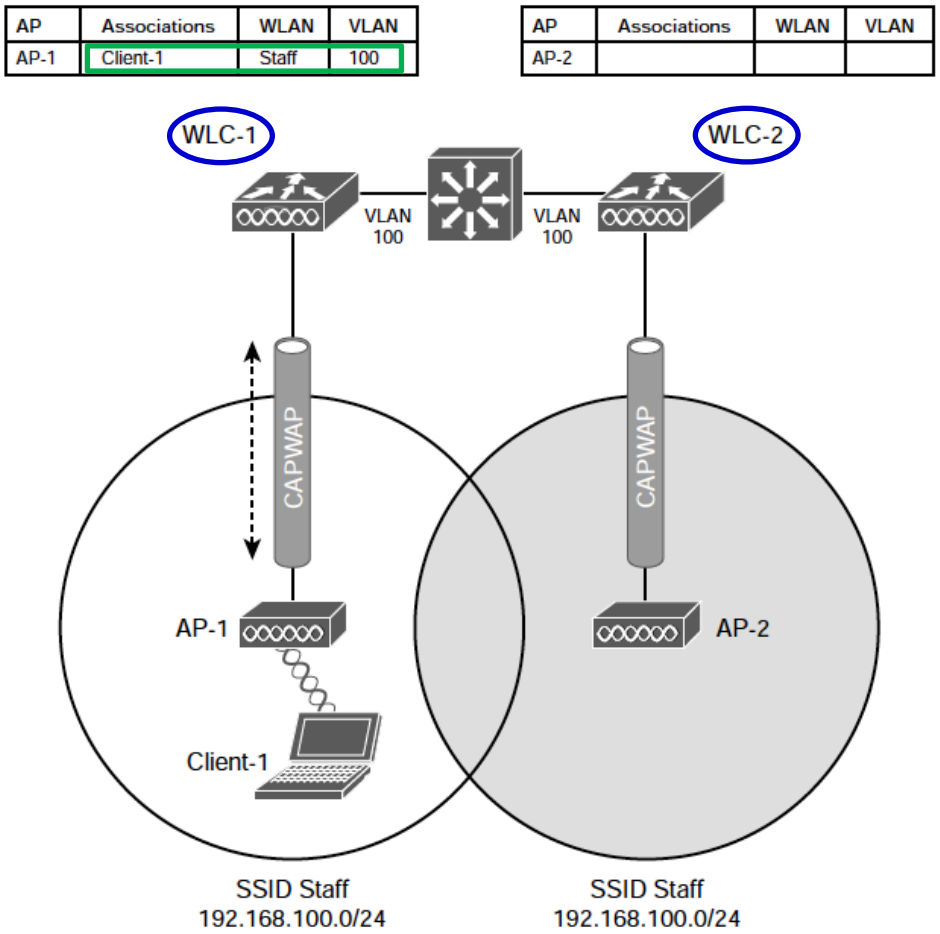

Figure 19-6 *Before an Intercontroller Roam*

After an Intercontroller Roam

Key Topic

AP	Associations	WLAN	VLAN

AP	Associations	WLAN	VLAN
AP-2	Client-1	Staff	100

When the client roams to a different AP, it can try to continue using its existing IP address or work with a DHCP server to either renew or request an address.

Figure 19-7 shows the client roaming to AP 2, where WLAN Staff is also bound to the same VLAN 100 and 192.168.100.0/24 subnet. Because the client has roamed between APs but stayed on the same VLAN and subnet, it has made a Layer 2 intercontroller roam.

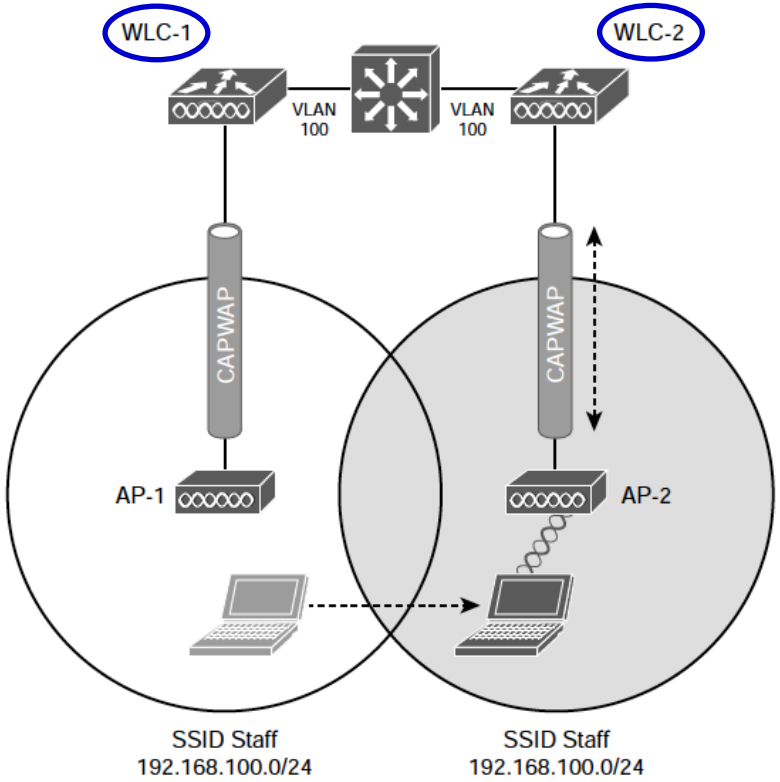

Figure 19-7 *After an Intercontroller Roam*

Before a Layer 3 Intercontroller Roam

When a client initiates an intercontroller roam, the two controllers involved can compare the VLAN numbers that are assigned to their respective WLAN interfaces. If the two VLAN IDs differ, the controllers arrange a Layer 3 roam that will allow the client to keep **using its IP address**.

Figure 19-8 illustrates a simple wireless network containing two APs and two controllers.

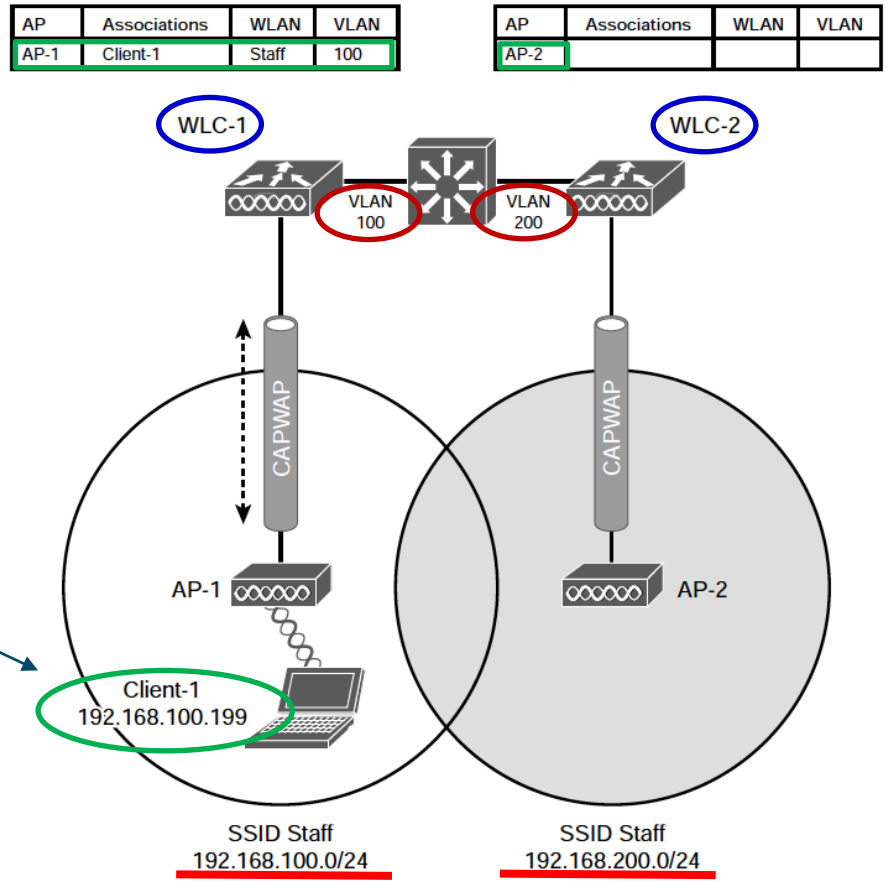

Figure 19-8 *Before a Layer 3 Intercontroller Roam*

After a Layer 3 Intercontroller Roam

A Layer 3 intercontroller roam consists of an **extra tunnel** that is built between the client's original controller **and** the controller it has roamed to. The tunnel carries data to and from the client as if it is still associated with the original controller and IP subnet.

Figure 19-9 shows the results of a Layer 3 roam.

The original controller (WLC 1) is called the **anchor** controller, and the controller with the roamed client is called the **foreign controller**.

Key Topic

AP	Associations	WLAN	VLAN
WLC-2	Client-1 (Mobile)	Staff	100

AP	Associations	WLAN	VLAN
AP-2	Client-1	Staff	

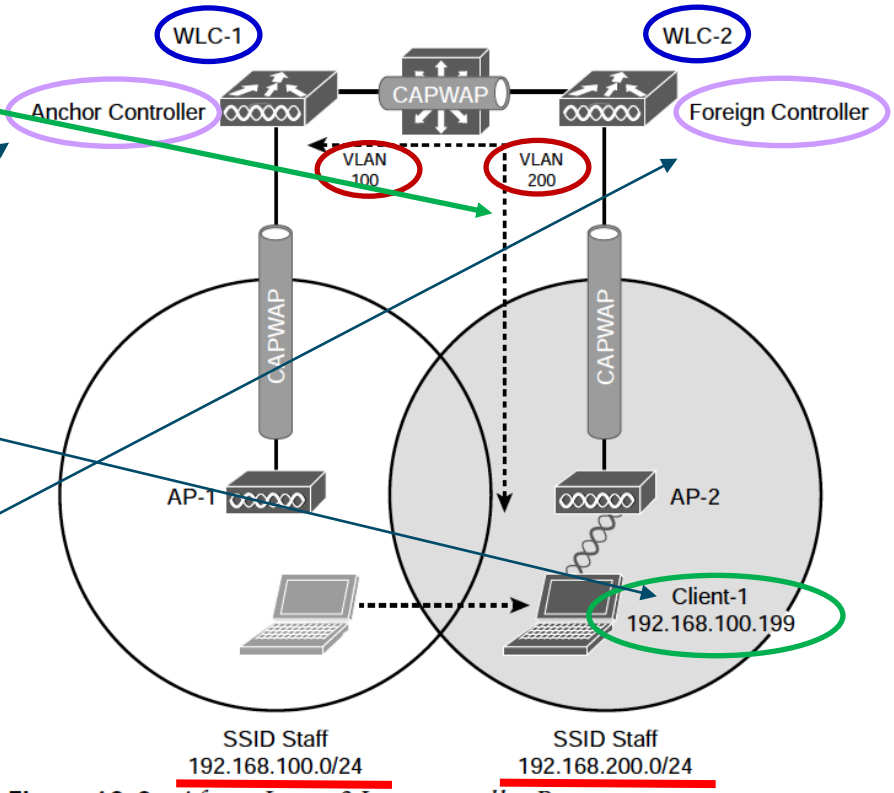

Figure 19-9 *After a Layer 3 Intercontroller Roam*

Scaling Mobility with Mobility Groups

Cisco controllers can be organized into mobility groups to facilitate intercontroller roaming.

If two centralized controllers are configured to belong to the same mobility group, clients can roam quickly between them.

If two controllers are assigned to different mobility groups, clients can still roam between them, but the roam is not very efficient. Credentials are not cached and shared, so clients must go through a full authentication during the roam.

Mobility Group Hierarchy

Key Topic

Mobility groups have an implied hierarchy, as shown in Figure 19-10.

Each controller maintains a mobility list that contains its own MAC address and the MAC addresses of other controllers. Each controller in the list is also assigned a mobility group name.

The mobility list gives a controller its view of the outside world; it knows of and trusts only the other controllers configured in the list.

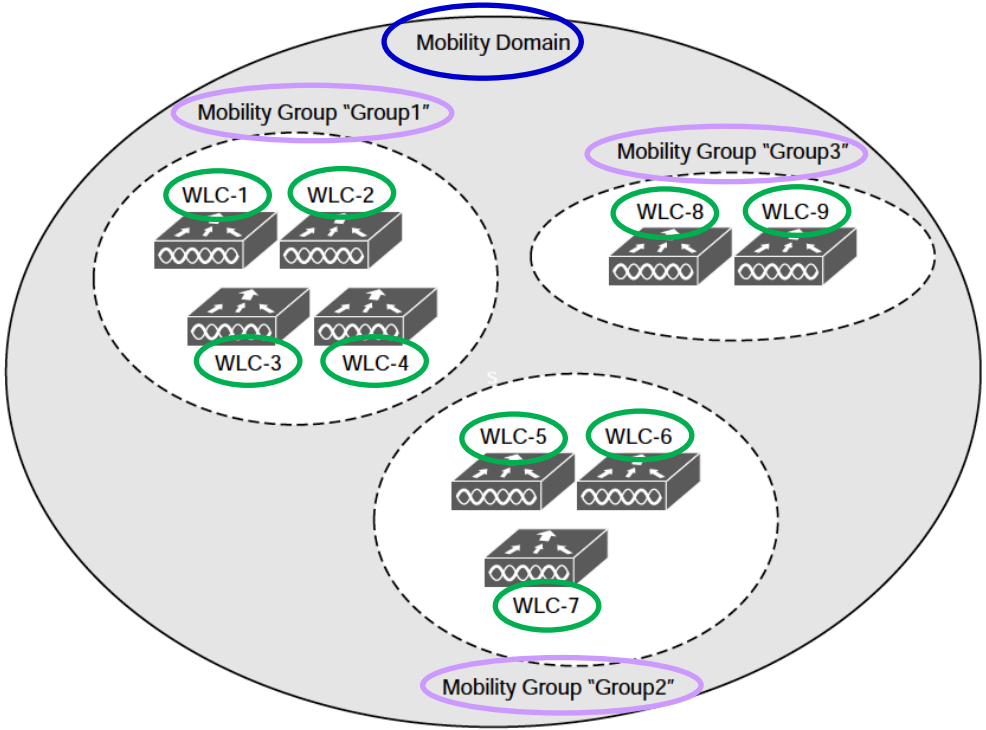

Figure 19-10 *Mobility Group Hierarchy*

Locating Devices in a Wireless Network

- Wireless networks are usually designed to provide coverage and connectivity in all areas where client devices are expected to be located. For example, a hospital building will likely have seamless wireless coverage on all floors and in all areas where users might go. Locating a user or device is important in several use cases, and a wireless network can be leveraged to provide that information.

Locating a Wireless Device with One AP or Three APs

A client's distance from an AP can be computed from its Received Signal Strength (RSS). If the distance is measured from a single AP only, it is difficult to determine where the client is situated in relation to the AP. Obtain the same measurement from three or more APs, then correlate the results and determine where they intersect.

Figure 19-11 illustrates the difference in determining a client's location with a single and multiple APs.

V zásadě klasická triangulace ☺

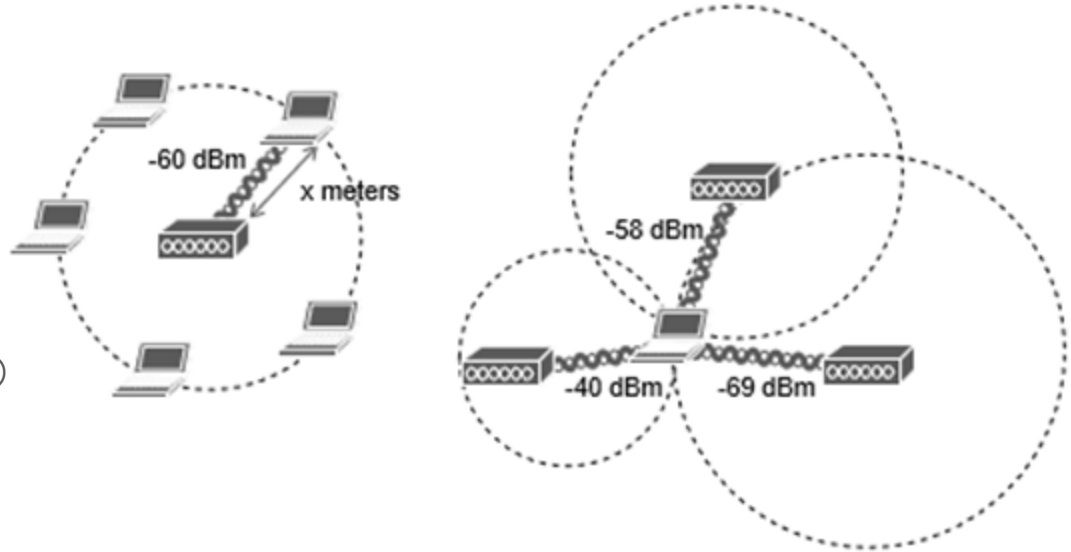

Figure 19-11 *Locating a Wireless Device with One AP (left) and Three APs (right)*

Real Time Location Data for Tracked Devices

The most intuitive way to interpret location data is to view devices on a map that represents the building and floor where they are located.

Figure 19-12 shows an example map of one floor of a building from Cisco DNA Spaces. The square icons represent AP locations, which were manually entered on the map. One device has been selected in the figure causing lines to be drawn to some of the APs that overheard the device.

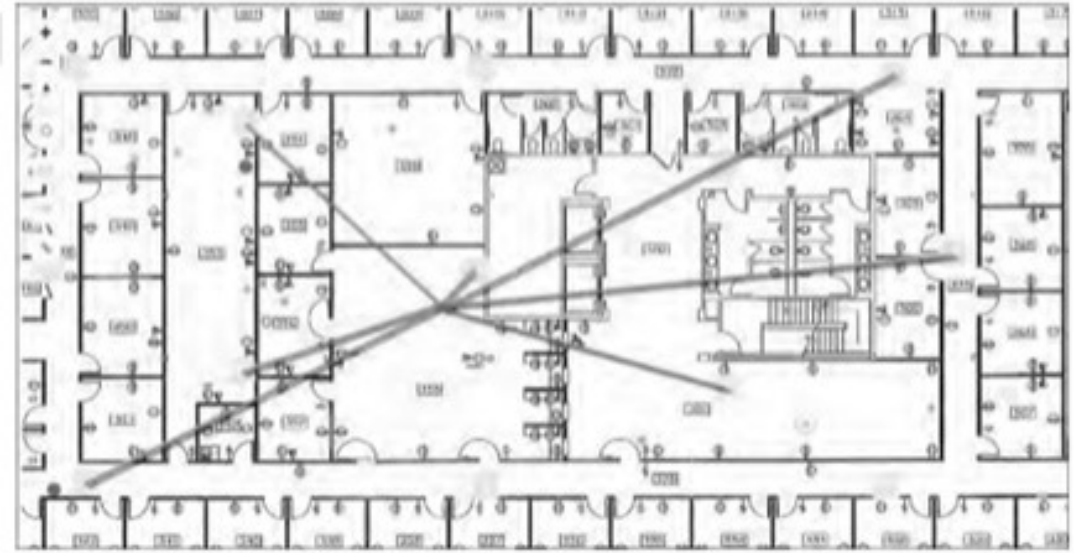

Figure 19-12 *An Example Map Showing Real Time Location Data for Tracked Devices*

Real Time Location for Other Tracked Devices

The same real-time location service also supports wireless devices that might never actually associate with an AP. For example, you might be interested in locating or tracking a potential customer's smartphone as he walks through a store. As long as Wi-Fi is enabled on the device, it will probably probe for available APs.

RFID tags are another type of device that can be attached to objects so that they can be tracked and located. Some RFID tags can actively join a wireless network to exchange data, while others are meant to simply "wake up" periodically to send 802.11 Probe Requests or multicast frames to announce their presence.

Another interesting use case is **locating rogue devices** and **sources of Wi-Fi interference**. Rogue devices will likely probe the network and can be discovered and located. Interference sources, such as cordless phones, wireless video cameras, and other transmitters, might not be compatible with the 802.11 standard at all.