Chapter 18: Wireless Infrastructure

Instructor Materials

CCNP Enterprise: Core Networking

Chapter 18 Content

This chapter covers the following content:

Wireless LAN Topologies - This section describes autonomous, cloud-based, centralized, embedded, and Mobility Express wireless architectures.

Pairing Lightweight APs and WLCs - This section explains the process that lightweight APs must go through to discover and bind to a wireless LAN controller.

Leveraging Antennas for Wireless Coverage - This section provides an overview of various antenna types and explains how each one alters the RF coverage over an area.

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.

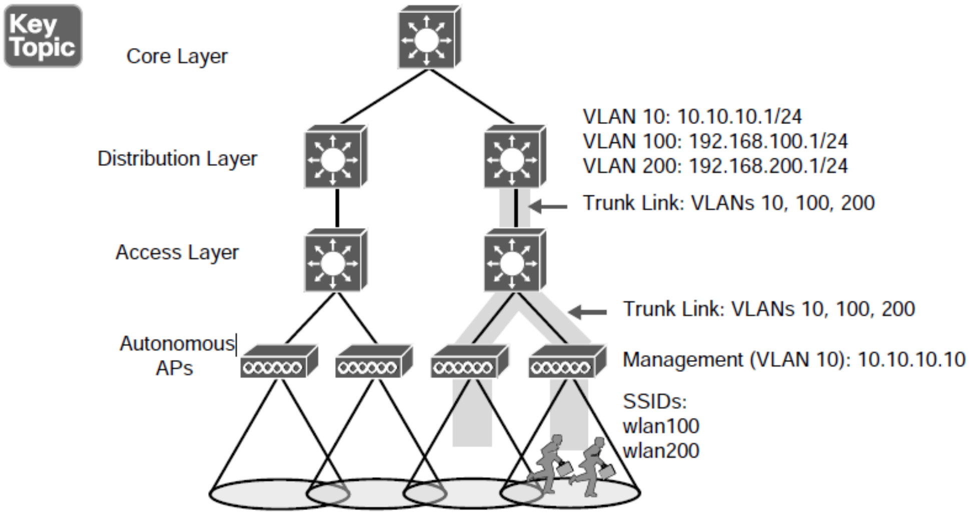

Figure 18-1 *Wireless Network Topology Using Autonomous APs*

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.

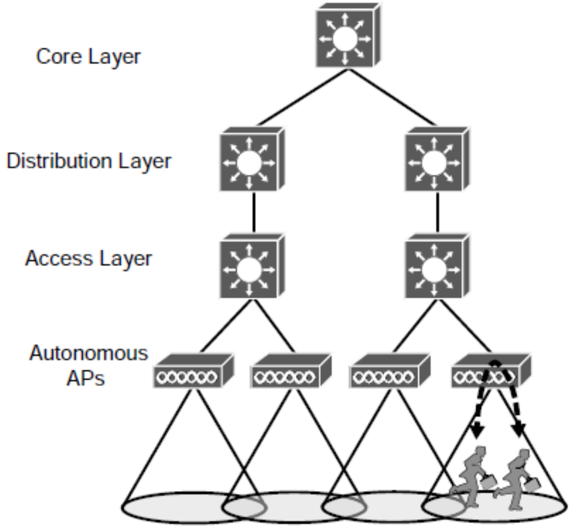

Figure 18-2 *Shortest Data Path Through an Autonomous AP Topology*

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

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.

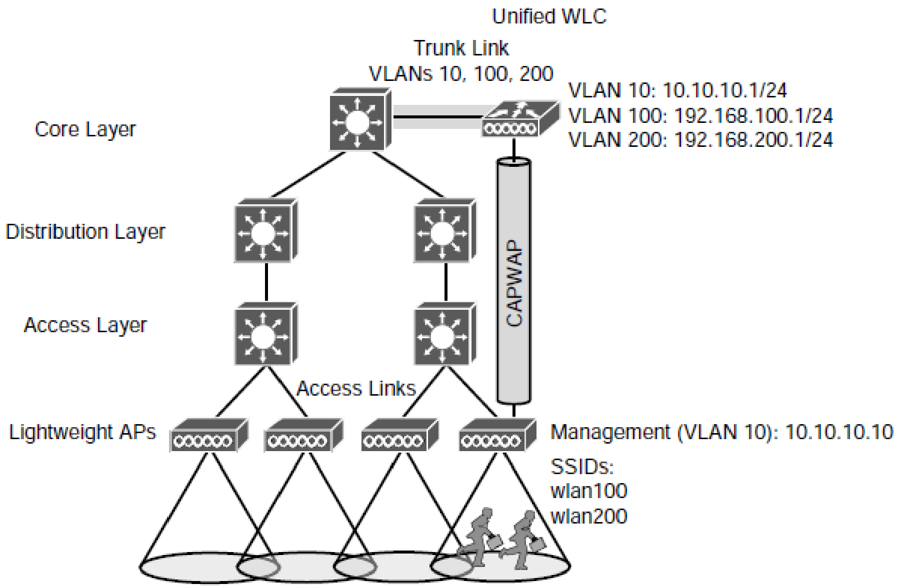

Figure 18-3 *WLC Location in a Centralized Wireless Network Topology*

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.

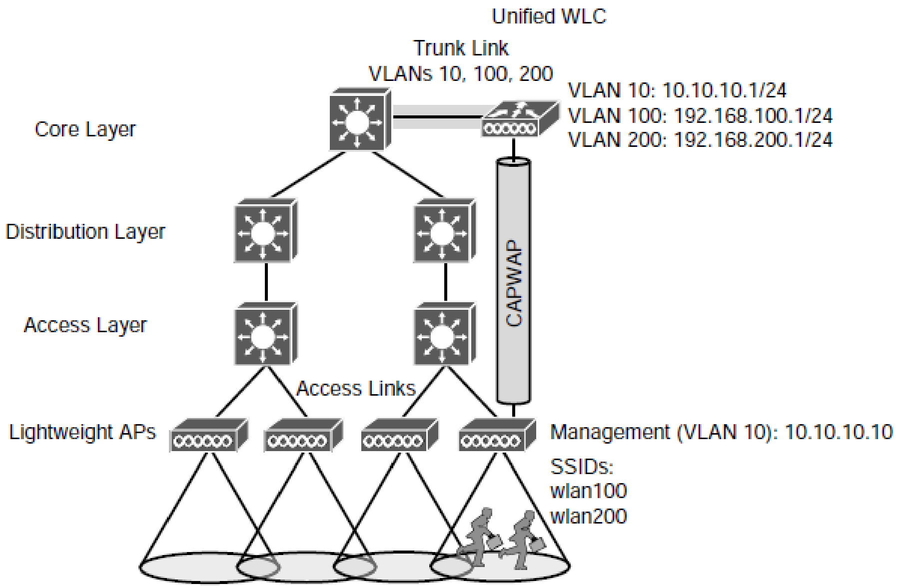

Figure 18-3 *WLC Location in a Centralized Wireless Network Topology*

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.

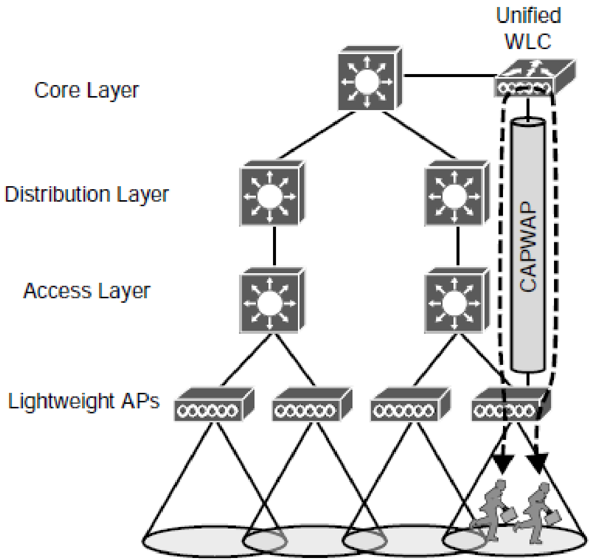

Figure 18-4 *Shortest Data Path Through a Unified Wireless Network Topology*

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.

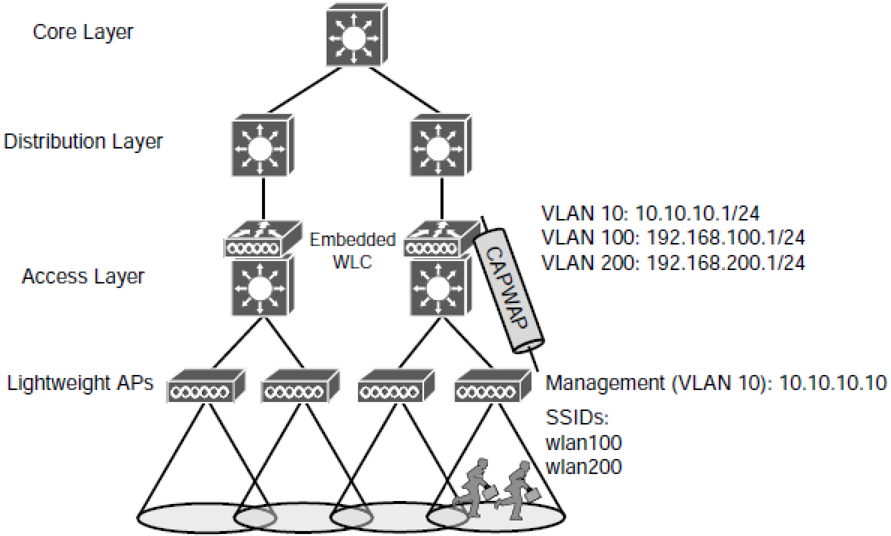

Figure 18-5 *WLC Location in an Embedded Wireless Network Topology*

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.

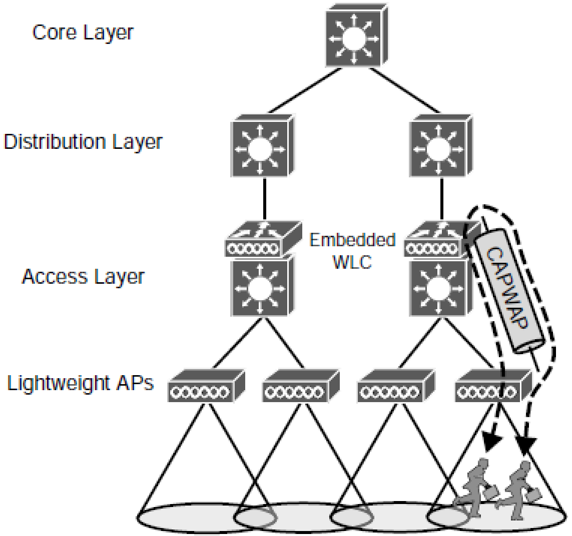

Figure 18-6 *The Shortest Data Path Through an Embedded Wireless Network Topology*

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.

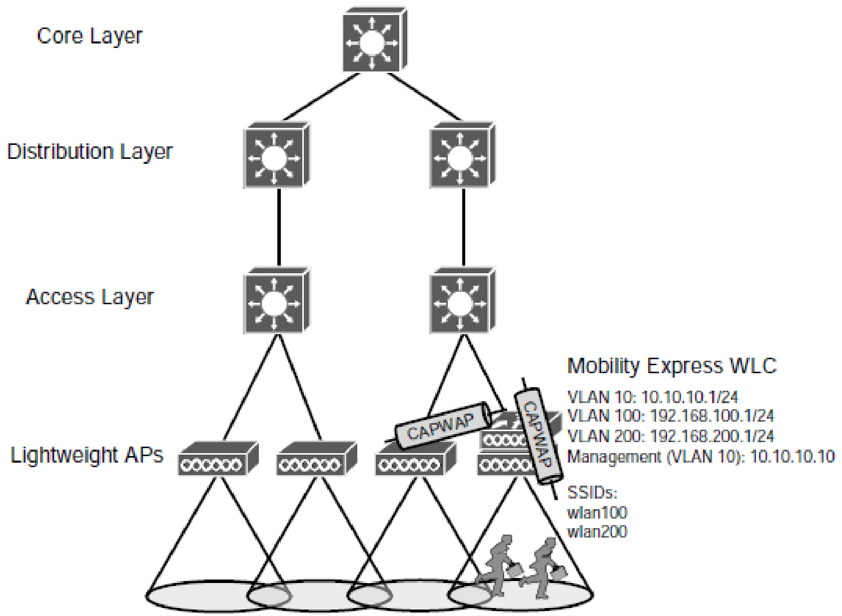

Figure 18-7 *WLC Location in a Mobility Express Wireless Network Topology*

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.

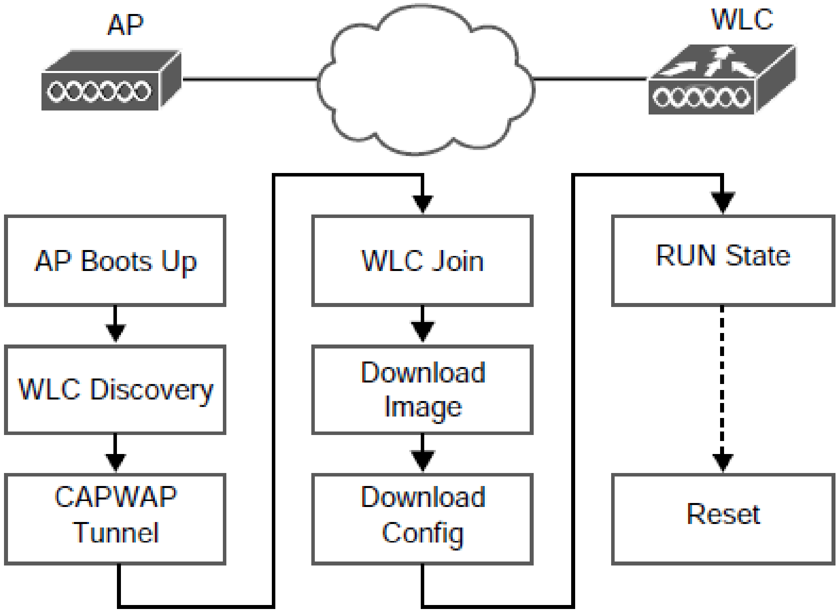

Figure 18-8 *State Machine of a Lightweight AP*

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.

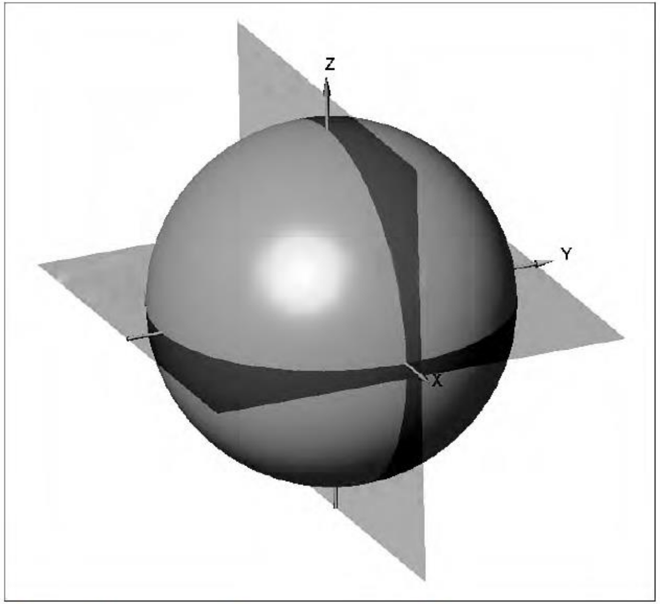

Figure 18-9 *Plotting the Radiation Pattern of an Isotropic Antenna*

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.

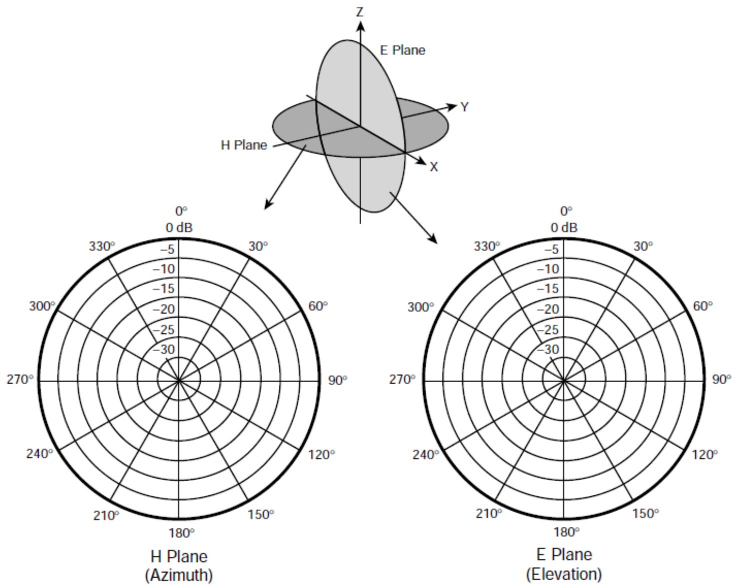

Figure 18-10 *Recording an Isotropic Antenna Pattern on E and H Polar Plots*

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.

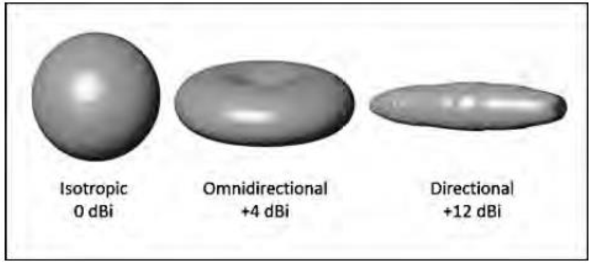

Figure 18-11 *Radiation Patterns for the Three Basic Antenna Types*

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.

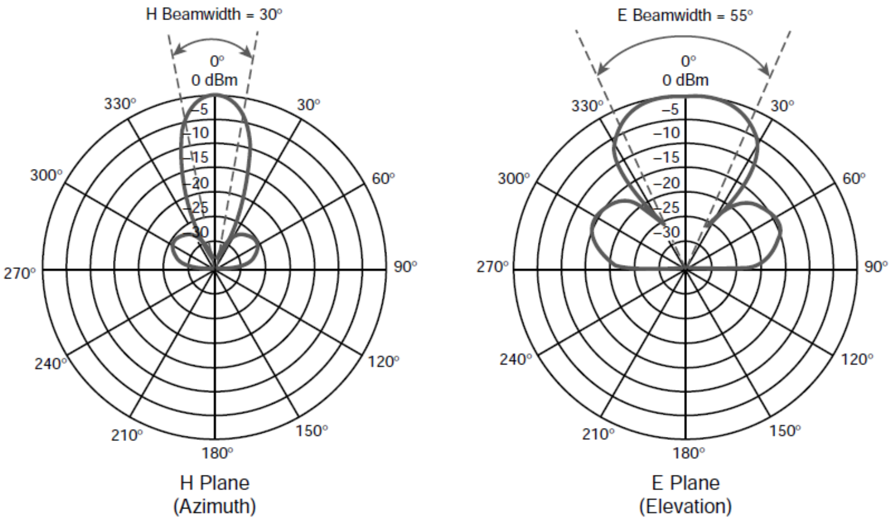

Figure 18-12 *Example of Antenna Beamwidth Measurement*

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.

Leveraging Antennas for Wireless Coverage

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.

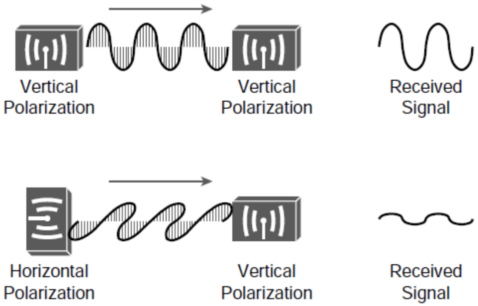

Figure 18-13 *Matching the Antenna Polarization Between Transmitter and Receiver*

Leveraging Antennas for Wireless Coverage
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.

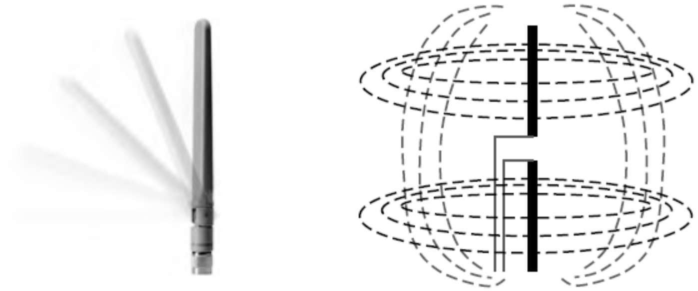

Figure 18-14 *Cisco Dipole Antenna*

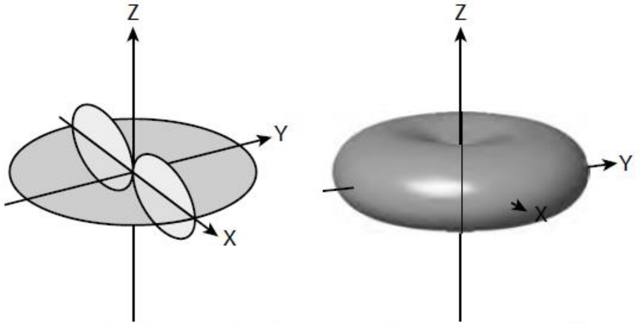

Figure 18-16 *Dipole Radiation Pattern in Three Dimensions*

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*

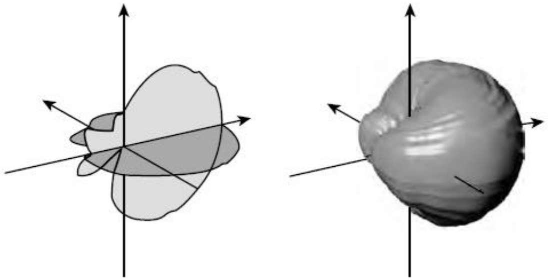

Figure 18-21 *Patch Antenna Radiation Pattern in Three Dimensions*

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*

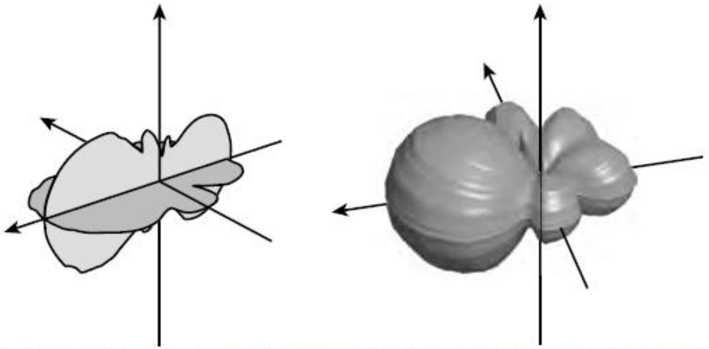

Figure 18-24 *Yagi Antenna Radiation Pattern in Three Dimensions*

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.

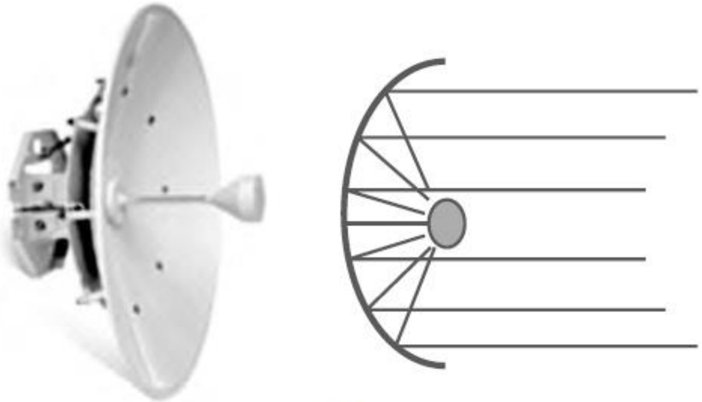

Figure 18-25 *Cisco Parabolic Dish Antenna*

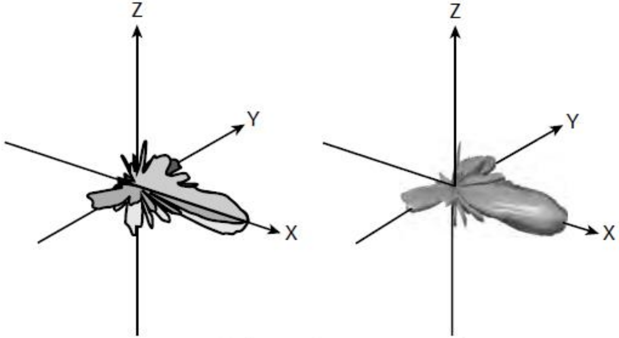

Figure 18-27 *Parabolic Dish Antenna Radiation Pattern in Three Dimensions*

Prepare for the Exam

Prepare for the Exam
Key Topics for Chapter 18

Description
Wireless Network Topology Using Autonomous APs
WLC Location in a Centralized Wireless Network Topology
WLC Location in an Embedded Wireless Network Topology
WLC Location in a Mobility Express Wireless Network Topology
AP controller discovery states
AP controller discovery steps
Cisco lightweight AP modes
Plotting the Radiation Pattern of an Isotropic Antenna

Prepare for the Exam

Key Terms for Chapter 18

Description	Description
Autonomous AP	E plane
Beamwidth	Embedded WLC deployment
CAPWAP	Gain
Centralized WLC deployment	H plane
Dipole	Integrated antenna
Directional Antenna	Lightweight AP

Prepare for the Exam

Key Terms for Chapter 18 (Cont.)

Description	Description
Local mode	Polarization
Mobility Express WLC deployment	Radiation pattern
Omnidirectional antenna	Split-MAC architecture
Parabolic dish antenna	Unified WLC deployment
Patch antenna	Wireless LAN Controller (WLC)
Polar plot	Yagi antenna