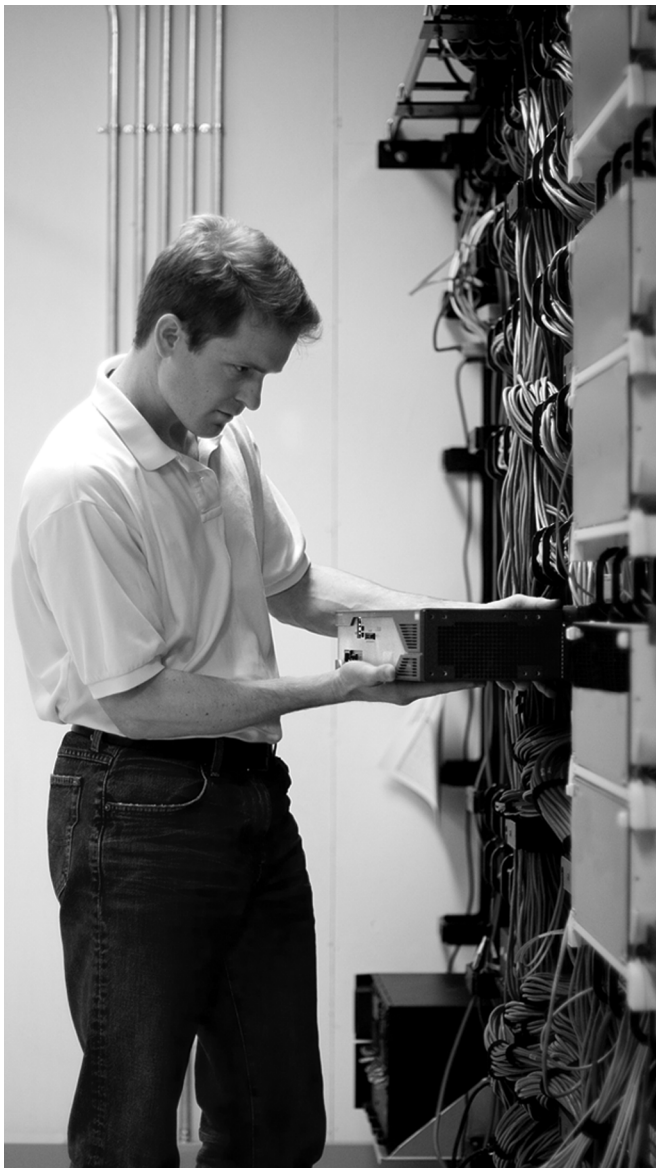

This chapter covers the following topics that you need to master for the CCNP BCMSN exam:

- **WLAN Security**—This section presents an overview of the various methods for protecting a WLAN. These methods can be used to authenticate potential wireless clients and users, as well as to secure the data passing over the wireless medium.
- **Wireless Client Operation**—This section explains the Cisco Compatible Extensions program and how it is used to find wireless hardware with compatible feature sets.
- **AP Association and Roaming**—This section covers the process that wireless clients and access points use to form associations or logical connections. As wireless clients become mobile, their associations can be moved to other access points. This forms the basis of client roaming.
- **Cell Layout and Channel Usage**—This section discusses the theory behind sizing and positioning access points so that they can work together to cover a large area. Part of this layout process is the assignment of RF channels and their distribution over the access point population.

CHAPTER 18

Wireless Architecture and Design

"Do I Know This Already?" Quiz

The purpose of the "Do I Know This Already?" quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz, derived from the major sections in the "Foundation Topics" portion of the chapter, helps you determine how to spend your limited study time.

Table 18-1 outlines the major topics discussed in this chapter and the "Do I Know This Already?" quiz questions that correspond to those topics.

Table 18-1 *"Do I Know This Already?" Foundation Topics Section-to-Question Mapping*

Foundation Topics Section	Questions Covered in This Section	Score
WLAN Security	1–5	
Wireless Client Operation	6	
AP Association and Roaming	7–10	
Cell Layout and Channel Usage	11–12	
Total Score		

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark this question wrong. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might give you a false sense of security.

1. If WPA or WPA2 is used as a wireless security suite, which one of the following represents a feature that is not offered in other wireless security methods?

   a. Virus mitigation

   b. Packet authentication

   c. Strong password policy

   d. Session time limits

2. If WPA is used, how often can the encryption keys be generated?
   a. One key per AP association
   b. One key per client session
   c. One key per packet
   d. One key per minute

3. TKIP is a protocol that is used with WPA to do which one of the following functions?
   a. Negotiate static WEP keys
   b. Generate per-packet encryption keys
   c. Authenticate a wireless user
   d. Authenticate wireless packets

4. Most wireless security methods use which one of the following as the authentication mechanism?
   a. 802.1D
   b. 802.11x
   c. RADIUS
   d. 802.1X

5. Which one of the following is a security feature that is unique to WPA2?
   a. AES encryption
   b. 3DES encryption
   c. WEP encryption
   d. Token-based encryption

6. Which one of the following can be used to verify feature compatibility between wireless devices?
   a. IEEE 802.11b
   b. IEEE 802.11e
   c. CCX
   d. CCO

7. Which one of the following determines when a wireless client will roam from one AP to another?
   a. The current AP has a weak signal from the client and asks it to roam
   b. The next AP overhears the client and asks it to roam
   c. The client's roaming algorithm reaches a threshold
   d. The client loses its IP address

8. Which one of the following is moved when a wireless client roams to a new AP?

   a. Association
   b. Certificate
   c. Beacon
   d. Channel

9. When a wireless client is actively roaming, which one of the following actions does it take?

   a. It listens for 802.11 beacons
   b. It listens for 802.11 Roam advertisements
   c. It sends an 802.11 Roam Request
   d. It sends an 802.11 Probe Request

10. Wireless client roaming from AP to AP normally occurs at what layer of the OSI model?

    a. Layer 1
    b. Layer 2
    c. Layer 3
    d. Layer 4

11. Which channels should be used across the 802.11b APs that are covering a floor of a building?

    a. 1, 2, 3
    b. 1, 3, 6
    c. 1, 3, 6, 11
    d. 1, 6, 11
    e. Any channel is fine

12. When you are designing the AP channel layout for an area, which one of the following is the most important consideration?

    a. The number of channels is conserved
    b. APs in different areas use different channels
    c. Adjacent APs use nonoverlapping channels
    d. Clients are grouped into common channels

The answers to the "Do I Know This Already?" quiz are found in Appendix A, "Answers to Chapter 'Do I Know This Already?' Quizzes and Q&A Sections." The suggested choices for your next step are as follows:

- **10 or less overall score**—Read the entire chapter. This includes the "Foundation Topics," "Foundation Summary," and "Q&A" sections.
- **11 or more overall score**—If you want more review on these topics, skip to the "Foundation Summary" section and then go to the "Q&A" section at the end of the chapter. Otherwise, move to Chapter 19, "Cisco Unified Wireless Network."

Foundation Topics

WLAN Security

As the central hub of a Basic Service Set (BSS), an AP effectively manages the WLAN for all clients within its range. Remember that all traffic going to or from a wireless client must go *through* the AP to reach other WLAN clients in the BSS or wired clients located elsewhere as illustrated in Figure 18-1. Clients cannot communicate directly with each other.

Figure 18-1 *An AP Serving as the Central Point of Contact in a WLAN*

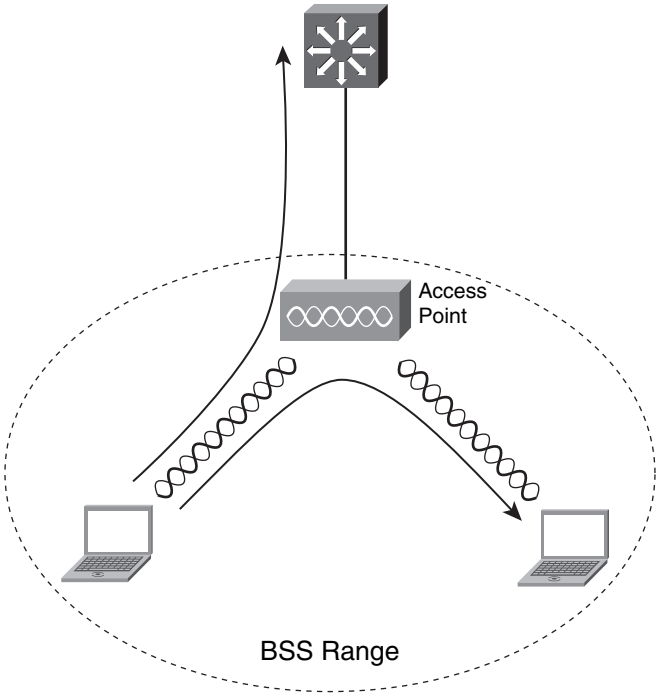

The AP is a natural place to implement various forms of security. For instance, an AP can control WLAN membership by authenticating clients. If a client fails to authenticate itself successfully, it won't be allowed to use the wireless network. As well, the AP and its clients can work together to secure the data that is passed between them. Otherwise, data sent over the air might be intercepted and exploited.

As a client brings up its wireless connection, it must find an AP that is reachable and that will approve its membership. The client must negotiate its membership and security measures in the following sequence, as shown in Figure 18-2:

1. Use an SSID that matches the AP
2. Authenticate with the AP
3. *(optional)* Use a packet encryption method (data privacy)
4. *(optional)* Use a packet authentication method (data integrity)
5. Build an association with the AP

Figure 18-2 *Basic Processes of Securing a Wireless LAN Connection*

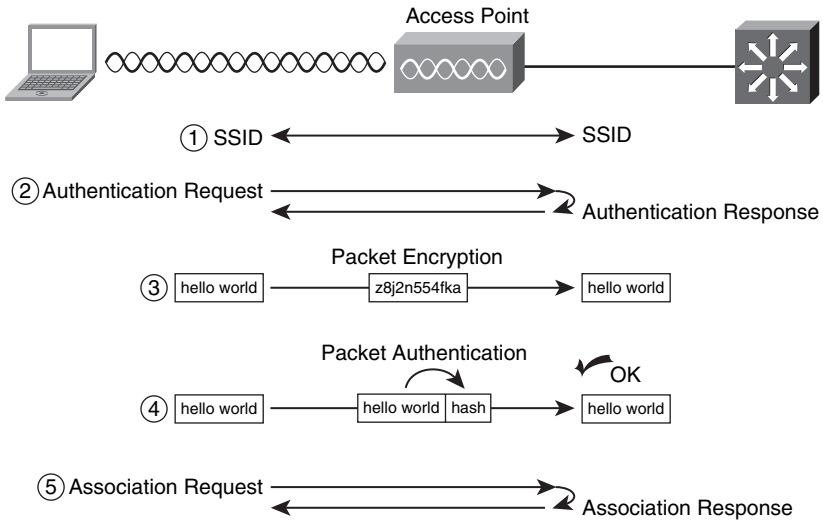

Chapter 17, "Wireless LAN Overview," discussed how the SSID string is used to match clients with the appropriate WLAN (and subsequent VLAN on the wired network). If a client's SSID is identical to the SSID used by an AP, the client can begin to communicate with the AP. The SSID is not meant to be used as any sort of security measure; its sole purpose is to break up a WLAN into logical groups of users.

Two basic concerns that 802.11 clients and APs must work out are authentication and encryption. Many different methods are available for authentication, encryption, and a combination of the two. The sections that follow briefly describe these methods.

Legacy Security

In 802.11 networks, clients can authenticate with an AP using one of the following methods:

- **Open authentication**—No authentication method is used; any client is offered open access to the AP.

- **Pre-shared key (PSK)**—The same secret key is statically defined on the client and the AP. If the keys match, the client is permitted to have access.

Notice that the authentication process in these two methods stops at the AP. In other words, the AP has enough information on its own to independently determine which clients can or can't have access. Open authentication and PSK are considered to be legacy methods because they are not scalable and are not necessarily secure.

Open authentication is usually the default, and offers no client screening whatsoever. Any client is permitted to join the network without presenting any credentials. In effect, the SSID is the only credential that is required! Although this makes life easier, it doesn't do much to control access to the WLAN. In addition, open authentication doesn't provide a means to encrypt data sent over the WLAN.

Pre-shared key authentication uses a long Wireless Equivalence Protocol (WEP) key that is stored on the client and the AP. When a client wants to join the WLAN, the AP presents it with a challenge phrase. The client must use the challenge phrase and the WEP key to compute a value that can be shared publicly. That value is sent back to the AP. The AP uses its own WEP key to compute a similar value; if the two values are identical, the client is authenticated.

When pre-shared key authentication (commonly called *static WEP keys*) is used, the WEP key also serves as an encryption key. As each packet is sent over the WLAN, its contents and the WEP key are fed into a cryptographic process. When the packet is received at the far end, the contents are unencrypted using the same WEP key.

Pre-shared key authentication is more secure than open authentication, but it has two shortcomings:

- It doesn't scale well because a long key string must be configured into every device.
- It isn't very secure.

As you might expect, a static key persists for a very long time, until someone manually reconfigures a new key. The longer a key remains in use, the longer malicious users can gather data derived from it and eventually reverse-engineer the key. It is commonly known that static WEP keys can be broken, so this method is not recommended.

EAP-Based Security Methods

Fortunately, wireless security has evolved to use other more robust methods. APs can use a variety of authentication methods that leverage external authentication and authorization servers and their user databases.

The Extensible Authentication Protocol (EAP) forms the basis for many wireless security methods—most of which have similar acronyms that rhyme, such as EAP, PEAP, and LEAP. EAP is defined in RFC 3748, and was originally designed to handle user authentication for PPP users.

Because it is extensible, it is well suited for a variety of security environments. RFC 4017 covers the EAP variants that are used in WLANs.

EAP has its history in PPP communication—not in wireless authentication. Chapter 15, "Securing Switch Access," described the IEEE 802.1x protocol as port-based authentication, or the means to authenticate users to use switch ports. Through 802.1x, users can authenticate even at Layer 2, before gaining further network connectivity. WLANs can leverage 802.1x as the means to implement EAP at Layer 2 for wireless clients.

In a wireless LAN, you can find some of the following security method names: LEAP, PEAP, EAP-TLS, and EAP-FAST. So many different methods exist that becoming confused about what they are and what they do is easy. Just remember that each one is based on EAP and uses a different type of credentials to authenticate wireless users.

Some of the EAP-based methods go beyond authentication by adding extra security features, as you will see as each method is discussed in the following sections.

LEAP

Cisco developed a protocol called Lightweight EAP (LEAP or EAP-Cisco) to address some shortcomings in 802.11 security. With LEAP, an AP uses an external Remote Authentication Dial-In User Server/Service (RADIUS) server to handle the actual client authentication. In fact, the AP and wireless client authenticate each other using a challenge and response exchange through the RADIUS server. Usernames and passwords are used as credentials.

LEAP also addresses wireless data privacy by assisting with WEP key assignment. A unique WEP key is dynamically generated by the RADIUS server for each wireless client. This process provides fresh encryption key material on a per-client basis, each time the client authenticates, and eliminates the need to manually configure static WEP keys altogether.

EAP-TLS

The EAP-TLS method, defined in RFC 2716, uses the Transport Layer Security (TLS) protocol to secure client authentication. TLS is based on Secure Socket Layer (SSL), which is commonly used in secure web browser sessions. EAP-TLS uses digital certificates as authentication credentials, which means that every AP and wireless client must have a certificate generated and signed by a common certificate authority (CA).

EAP-TLS also addresses wireless data privacy by generating WEP keys automatically, each time the authentication server forces the client to reauthenticate. The TLS session key, unique to each wireless client that is authenticating, is used to derive a unique WEP key. The WEP key is then used to encrypt the wireless data.

PEAP

Protected EAP (PEAP or EAP-PEAP) is similar to EAP-TLS in that a TLS session is used to secure the authentication. PEAP requires a digital certificate only on the authentication server so that the server itself can be authenticated to the client. The wireless clients are authenticated using Microsoft Challenge Handshake Authentication Protocol version 2 (MSCHAPv2).

As with EAP-TLS, the TLS session key is used to derive a WEP key for encrypting the wireless data stream. The keys change periodically as the authentication server forces the client to reauthenticate.

EAP-FAST

EAP Flexible Authentication via Secure Tunneling (EAP-FAST) is a wireless security method developed by Cisco. EAP-FAST is not named for its speed; rather, it is named for its flexibility to reduce the administrative complexity. Clients aren't required to use digital certificates, and they aren't required to follow strict or strong password policies.

EAP-FAST works by building a secure tunnel between the client and the authentication server. A Protected Access Credential (PAC) is used as the only client credential to build the tunnel. The PAC can be assigned from a PAC server or it can be created dynamically during a phase of EAP-FAST negotiations. Once the tunnel is built, the client is authenticated using familiar username and password credentials.

EAP-FAST can derive a WEP key dynamically so that the wireless data stream can be encrypted.

WPA

The IEEE 802.11i standard focuses on addressing all aspects of wireless security—even beyond client authentication and data privacy using WEP keys. As the 802.11i standard was being developed, wireless LAN vendors have moved ahead to implement as many of its features as possible. As a result, the Wi-Fi Alliance developed *Wi-Fi Protected Access (WPA)* based on some of the 802.11 draft components.

WPA offers the following wireless LAN security measures:

- Client authentication using 802.1x or a pre-shared key
- Mutual client-server authentication
- Data privacy using Temporal Key Integrity Protocol (TKIP)
- Data integrity using Message Integrity Check (MIC)

TKIP leverages existing WEP encryption hardware that is embedded in wireless clients and APs. The WEP encryption process remains the same, but the WEP keys are generated much more frequently than the periodic reauthentications that occur with EAP-based authentication methods.

In fact, TKIP generates new WEP keys on a *per-packet* basis! An initial key is built as a client authenticates (or reauthenticates) with the EAP-based method. That key is formed by mixing the MAC address of the transmitter (the client or the AP) with a sequence number. Each time a packet is sent, the WEP key is incrementally updated. Once the client is forced to reauthenticate, an entirely new WEP key is built and the per-packet process repeats.

WPA can use a pre-shared key for authentication if external authentication servers aren't used or required. In that case, the pre-shared key is used only during the mutual authentication between the client and the AP. Data privacy or encryption doesn't use that pre-shared key at all. Instead, TKIP takes care of the rapid encryption key rotation for WEP encryption.

The MIC process is used to generate a "fingerprint" for each packet sent over the wireless network. If the fingerprint is made just before the packet is sent, the same fingerprint should match the packet contents once the packet is received. Why bother fingerprinting packets in the first place? When packets are sent over the air, they can be intercepted, modified, and re-sent—something that should never be allowed to happen. Fingerprinting is a way to protect the integrity of the data as it travels across a network.

For each packet, MIC generates a hash code (key), or a complex calculation that can only be generated in one direction. The MIC key uses the original unencrypted packet contents and the source and destination MAC addresses in its calculation, so that these values can't be tampered with along the way. As shown in Figure 18-2 under "Packet Authentication," the MIC hash key is added to the original packet so that the receiving end can examine the key and detect any tampering.

WPA2

Wi-Fi Protected Access version 2 (WPA2) is based on the final 802.11i standard. WPA2 goes several steps beyond WPA with its security measures.

For data encryption, the Advanced Encryption Standard (AES) is used. AES is a robust and scalable method that has been adopted by the National Institute of Standards and Technology (NIST, www.nist.gov) for use in the U.S. government organizations. TKIP is still supported for data encryption, for backward compatibility with WPA.

With WPA and other EAP-based authentication methods, a wireless client has to authenticate at each AP it visits. If a client is mobile, moving from AP to AP, the continuing authentication process can become cumbersome. WPA2 solves this problem by using proactive key caching (PKC). A client authenticates just once, at the first AP it encounters. As long as other APs visited support WPA2 and are configured as one logical group, the cached authentication and keys are passed automatically.

Wireless Client Operation

Wireless devices can be purchased from a variety of vendors, each with its own set of features and requirements. As well, wireless clients can exist as internal or external adapters installed in PC platforms. They can also be embedded in other devices such as cell phones, wireless phones, PDAs, medical devices, and tags used for location tracking. These are usually called *application-specific devices (ASDs)*.

If you use Cisco APs in your network, knowing whether each wireless device is indeed compatible with the features you plan to use would be nice. Cisco has developed the Cisco Compatible Extensions (CCX) program to address this need. Before a device can be CCX-compatible, it must be fully tested and verified to be compatible, whichis especially handy when a pre-standard feature needs to be used, and you have no other guarantee that various vendors have implemented the feature in the same way.

As wireless LAN features have been introduced over time, the CCX program has evolved to include them. CCX is broken down into different versions, with each higher version containing all the features listed in lower versions. At press time, CCX version 4 was the most recent.

Table 18-2 breaks down the basic groups of features of the various CCX versions.

Table 18-2 *CCX Features*

CCX Version	Features Covered
CCXv1	Basic 802.11 and Wi-Fi compatibility 802.1X authentication for LEAP Multiple SSID use
CCXv2	WPA 802.1X authentication for PEAP Fast roaming with CCKM RF scanning for WLAN site survey and interference monitoring
CCXv3	WPA2, including AES encryption 802.1X authentication for EAP-FAST Wi-Fi Multimedia (WMM) as part of the 802.11e QoS standard
CCXv4	Cisco Network Admission Control (NAC) Call admission control for Voice over IP (VoIP) Reporting VoIP metrics Enhanced roaming 802.11 location tag functionality (radio frequency identification [RFID])

NOTE You can find the latest CCX program details in the "Cisco Compatible Extensions Program for Wireless LAN (WLAN) Client Devices" document, at www.cisco.com/web/partners/pr46/pr147/partners_pgm_concept_home.html. The most current information on CCX versions can be found in "Cisco Compatible Extensions: Versions and Features" at www.cisco.com/web/partners/pr46/pr147/program_additional_information_new_release_features.html.

AP Association and Roaming

When a wireless client is associated with an AP, all data going to and from the client must pass through that AP. Recall from Chapter 17 that a client forms an association by sending an association request message to the AP. If the client is compatible with the WLAN by having the correct SSID, supporting the same data rates, and authenticating correctly, the AP responds with an association reply.

An association is maintained with the AP as long as the client stays within range of the AP. Consider the AP cell shown in Figure 18-3. As long as the client stays within points A and B, it is able to receive the AP's signal at an acceptable level. As soon as the client goes outside the cell range at point C, the signal strength falls below the acceptable threshold and the client loses the association.

Figure 18-3 *A Mobile Client Moves Within an AP Cell*

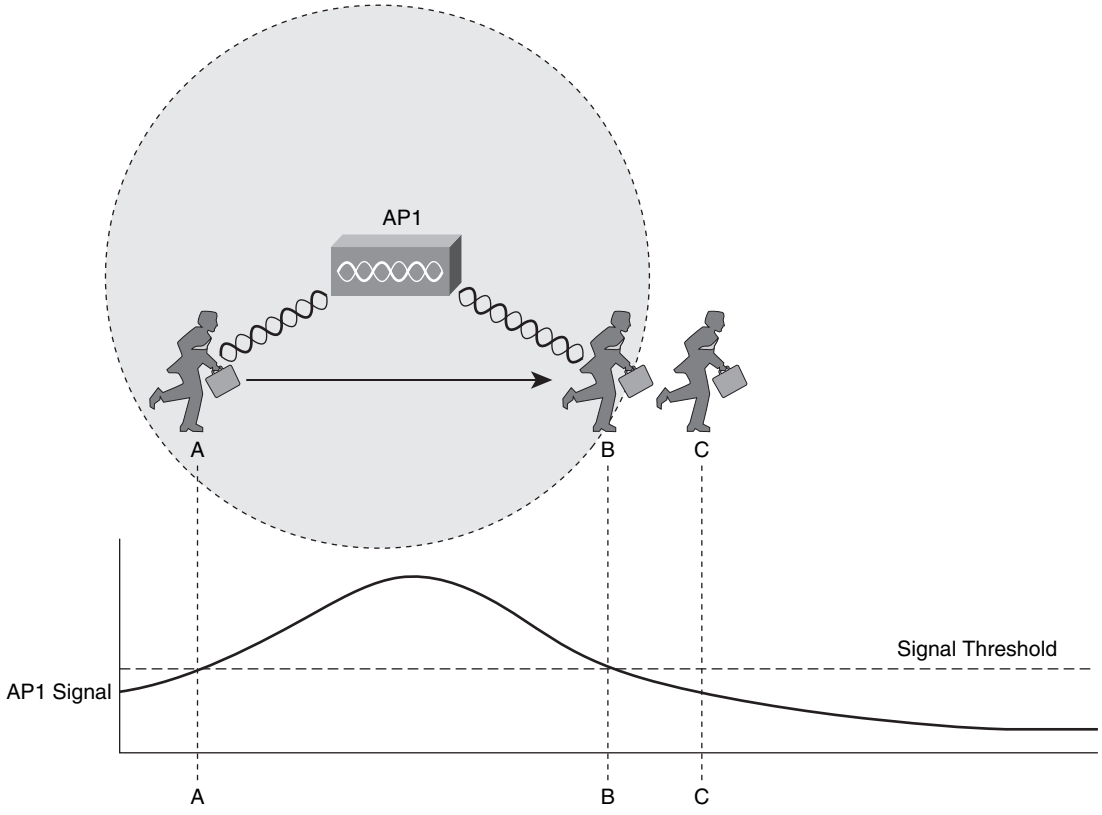

Other APs can be added so that the client can move within a larger area. However, the APs must be carefully deployed to allow the client to roam from AP to AP. *Roaming* is the process of moving an association from one AP to the next so that the wireless connection is maintained as the client moves.

In Figure 18-4, two APs are located side by side, each using the same channel. Building a large coverage area using a single channel might seem intuitive, but it turns out to be a bad idea because the client isn't able to decide when it has roamed away from one AP into the cell of another.

Figure 18-4 *Pitfalls of Reusing Channels in Adjacent APs*

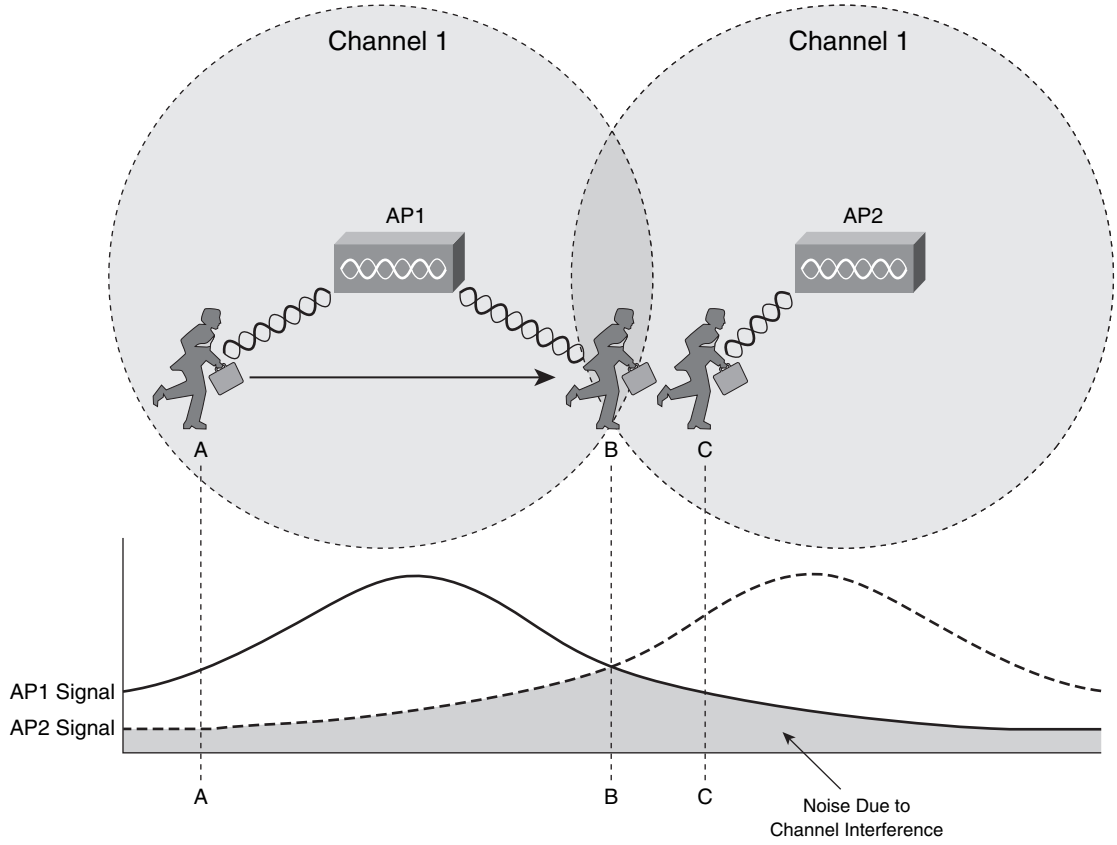

Remember that the signal from an AP doesn't actually stop at the edge of the cell—rather, it continues to propagate as it eventually dies off. This is shown by the signal strength graph of each AP. The client is able to form an association with AP1 at point A. Even at that location, some portion of AP2's signal can be received. Because AP2 is using the same channel as AP1, the two APs essentially interfere with each other.

Ideally, when the client in Figure 18-4 moves to location B, it should begin to anticipate the need to roam or transfer its association from AP1 to AP2. With channel interference from the two APs,

it might never be able to roam cleanly. In fact, the client might never be able to operate cleanly in either cell.

The Roaming Process

What enables a client to roam in the first place? First, adjacent APs *must* be configured to use different nonoverlapping channels. For example, APs operating under 802.11b or 802.11g must use only channels 1, 6, and 11. An AP using channel 1 must not be adjacent to other APs using channel 1. This ensures that clients will be able to receive signals from a nearby AP without interference from other APs.

The roaming process is driven entirely by the wireless client driver—not by the AP. The client can take two approaches to decide when to roam:

- The client can proactively search for other adjacent APs *before* it experiences the need to roam.
- The client can search for adjacent APs *after* it realizes that it needs to roam.

Wireless clients decide that it's time to roam based on a variety of conditions. The 802.11 standards don't address this issue at all, so roaming algorithms are vendor-specific. As well, the roaming algorithms are usually "secret recipes" so that the exact thresholds and conditions are hidden from view.

Some of the ingredients in the roaming algorithm are signal strength, signal quality, a count of missed AP beacons, errors due to collisions or interference, and so on. These items are usually logical choices because they indicate the overall quality of a connection.

Because different clients use different thresholds, some will try to roam earlier than others at a given location within a cell. Some clients will tend to "latch on" to an existing association until the AP can hardly be heard, whereas others will attempt to roam whenever a better AP can be reached. In other words, don't worry too much about what controls the roaming algorithm. Rather, just be familiar with the roaming process.

Once a client decides it should roam, it must begin to search for a new potential AP. It does so by scanning the other channels to find other active APs. The client can take two approaches to the scanning process:

- **Passive scanning**—The client takes time to scan other channels, but only listens for 802.11 beacons from available APs.
- **Active scanning**—The client takes time to scan other channels, but sends 802.11 Probe Request frames to query available APs.

When a client passively scans, it has to only wait to receive beacons. Therefore, it is well suited for low-power and embedded wireless clients. Active scanning puts the client in control because it must send probes and wait to receive probe replies. Active scanning usually results in more efficient roaming than passive scanning because APs can be queried and identified on-demand.

In Figure 18-5, two APs have been correctly configured with nonoverlapping channels 1 and 6. The two AP signal strengths are also shown as a graph corresponding to the client's location. At location A, the client has a clear signal from AP1, so it maintains an association with that AP.

Figure 18-5 *A Client Roaming Between Two APs*

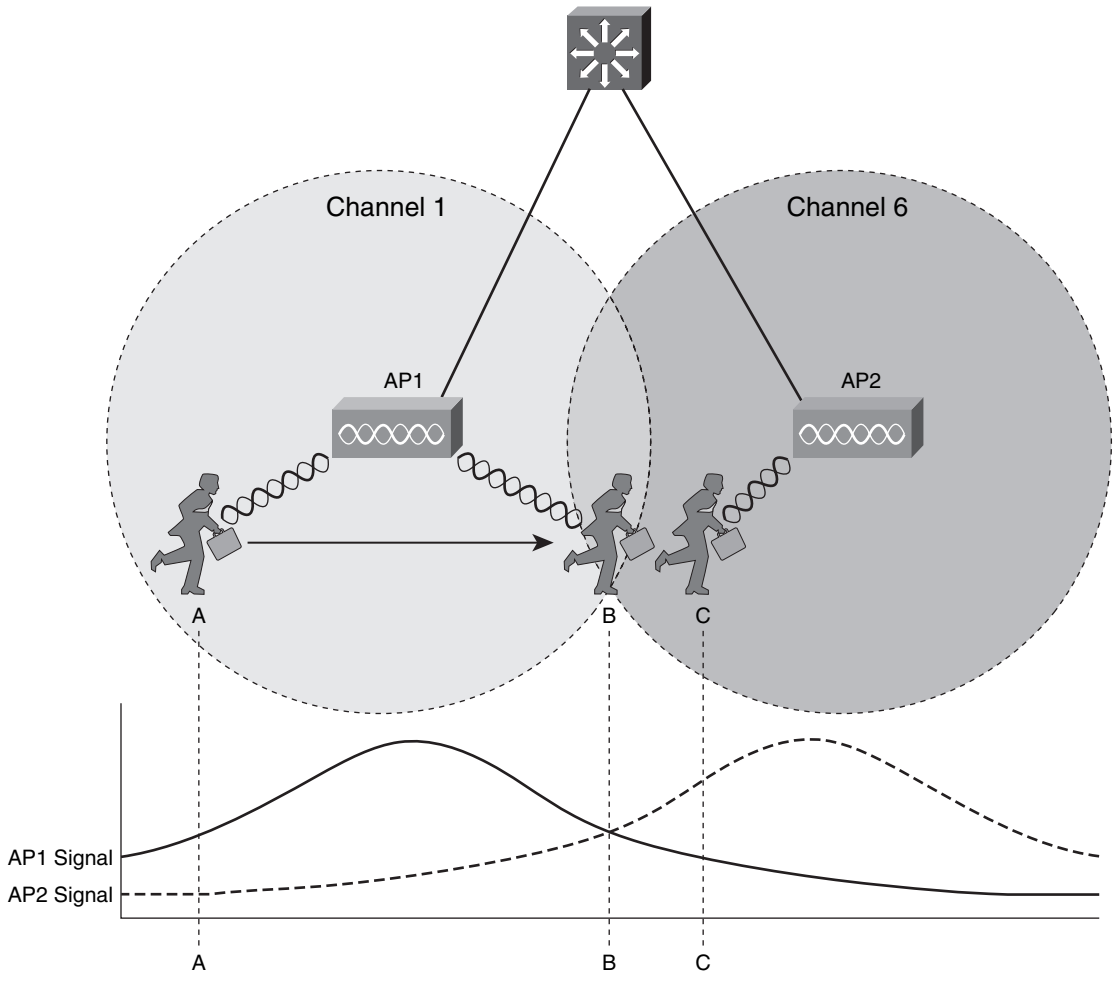

As the client moves toward location B, it decides that AP1's signal is no longer optimal. Somewhere along the way, the client begins to seek out a better AP where it can move its association. A wireless client does this in a two-step fashion:

Step 1 The client sends 802.11 *probe request* management frames to any listening AP.

Step 2 Any listening AP answers the client with 802.11 *probe response* frames, advertising the AP's existence.

The client doesn't know what channel is used on the next AP it encounters, so it must send the probes over every possible channel. Therefore, the client must take time to tune its radio away from the current AP's channel so it can scan other channels and send probes.

You might think of this as someone watching television. As the current program gets boring or nears its end, the viewer begins to "channel surf" and scans other channels for a better program.

One thing to keep in mind: while the viewer is scanning channels, he cannot keep watching the original program. Some of that program will be missed. This is also true of wireless clients. While a radio is scanning other channels, packets arriving on the original channel will be dropped because they can't be received. Therefore, a trade-off exists between staying available on a single channel and attempting to roam to other APs.

Returning to Figure 18-5, when the client nears location B, it sends 802.11 probe request frames on a variety of channels. When AP2 receives a probe request on channel 6, it replies with a probe reply on channel 6. After the client is satisfied with any probe replies it receives, it evaluates them to see which AP offers the most potential for a new association.

Now the client must roam and actually move its association. Notice in Figure 18-5 that the client is still associated with AP1 at location B, even though it might be able to receive AP2 as good or better.

First, the existing association must be dropped because a client is only permitted to associate with one AP at a time. The client sends an 802.11 disassociation message to AP1 over channel 1—the channel used by AP1. Then the client is free to send an association request to AP2 over channel 6, which is followed by an association response from AP2.

Roaming Implications

As Figure 18-5 hinted, adjacent APs are connected by a switched network and a single, common VLAN. Therefore, native 802.11 roaming between APs really takes place at Layer 2. You can think of this as if the wired connection for a client PC is moved from access layer switch to access layer switch—all within the same VLAN.

The implication here is that the client's IP address stays the same, even while roaming. This is handy because the client doesn't have to spend time acquiring a new IP address when it associates with a different AP.

During the roaming process, the client must release one association before negotiating the next association. There is a brief time when the client has no association with any AP. This is actual dead time when the client isn't able to send or receive data. However, the goal for Layer 2 roaming is to keep this dead time to a minimum so that delay-sensitive applications aren't adversely affected.

At some point, once the WLAN reaches a large size, it is better to start over with a new IP subnet and VLAN. From the earlier chapters in this book, you should recall that large campus networks should be broken down into switch blocks so that there aren't any end-to-end or campus VLANs. This is also important with WLANs, as they are really just an extension of the switched infrastructure.

If the WLAN is broken up into multiple VLANs and subnets, wireless clients might have to cross Layer 3 boundaries when they roam. At those locations, the client IP addresses will change from one AP to another. This involves more than simple 802.11 probes and association requests—it also requires additional dead time while the client requests and receives a new IP address.

Layer 3 roaming is not native to standard APs. It requires the leverage of other tools that can be overlaid on the 802.11 network. This problem can be solved with the wireless infrastructure that is described in Chapter 19, "Cisco Unified Wireless Network."

Cell Layout and Channel Usage

The previous section laid the foundation for roaming by describing movement between two AP cells. Most scenarios require more than two APs to cover the appropriate area within a building. Therefore, you need to consider the layout and configuration of more and more APs to scale the design to fit your wireless environment.

For example, to cover the entire area of a warehouse or one floor of a building, APs must be placed at regular intervals throughout that space. A site survey is a vital step toward deciding on AP placement, as actual live measurements are taken with an AP staged at various points in the actual space.

The two basic goals when designing a WLAN are

- Sizing the AP cells
- Selecting channels for the AP cells

The sections that follow describe these goals.

Sizing AP Cells

The size of AP cells determines the number of APs that must be purchased and deployed to cover an area; however, your design should not be driven by the cost alone. AP cell size can also affect the performance of the APs as clients move around or gather in one place.

Remember that a WLAN is a shared medium. Within a single AP cell, all the clients associated with that AP must share the bandwidth and contend for access. If the cell is large, a large number of clients could potentially gather and use that AP. If the cell size is reduced, the number of simultaneous clients can also be reduced.

TIP No clear rule of thumb exists for sizing AP cells for a specific number of clients. As with switched networks, the limiting factor is really the type of applications the clients will use, as well as the volume of data moving over the medium at any given time.

As a very loose guideline, you can consider the maximum peak throughput of a wireless cell divided by the number of simultaneous clients to get a maximum data rate per user. Factoring in the overhead of 802.11 encapsulation and bandwidth contention, 802.11b generally offers up to 6.8 Mbps through an AP, whereas 802.11g and 802.11a offer up to 32 Mbps.

This means, for example, in an 802.11b cell with 25 clients, each client would have a maximum throughput of 6.8 Mbps / 25, or 272 kbps. In an 802.11a or 802.11g cell, those same 25 users would have 32 Mbps / 25, or 1.28 Mbps.

You should also keep in mind that large cells can allow clients to step their data rates down as they move farther away from the APs. For example, when an 802.11b client is near an AP, it can use the highest data rate (11 Mbps). As the client moves out away from the AP, the data rate can be reduced to 5.5, 2, and finally 1 Mbps. You might want your clients to use only the highest data rates in a cell, which can be accomplished by reducing the cell size.

Generally, the AP cell size is driven by the AP's transmit power. Higher power equates to greater range, so the power must be adjusted so that the AP's signal doesn't propagate into nearby AP cells operating on the same channel.

TIP For more detailed information about AP cell size and wireless LAN site surveys, refer to the Cisco Airespace Installation, Administration, and Maintenance (CAIAM) or Cisco Aironet Wireless Site Survey (CAWSS) course.

Once the AP cells have been sized and pinpointed, clients should be able to associate and roam at any location within the coverage area. If one AP should fail, the area it originally covered will be left silent.

Naturally, this "hole" in the coverage can be fixed by simply replacing the failed AP—assuming that you could discover the failed AP radio in the first place. In the meantime, you could also configure adjacent APs to increase their output power to expand their coverage area over the hole. However, tweaking the AP power is a tricky task that can affect many other AP cells as well.

WLAN Channel Layout

To minimize channel overlap and interference, AP cells should be designed so that adjacent APs use different channels. With 802.11b and 802.11g, you are limited to using channels 1, 6, and 11. The cells could be laid out in a regular, alternating pattern, as Figure 18-6 illustrates.

Figure 18-6 *Holes in an Alternating Channel Pattern in 802.11b/g*

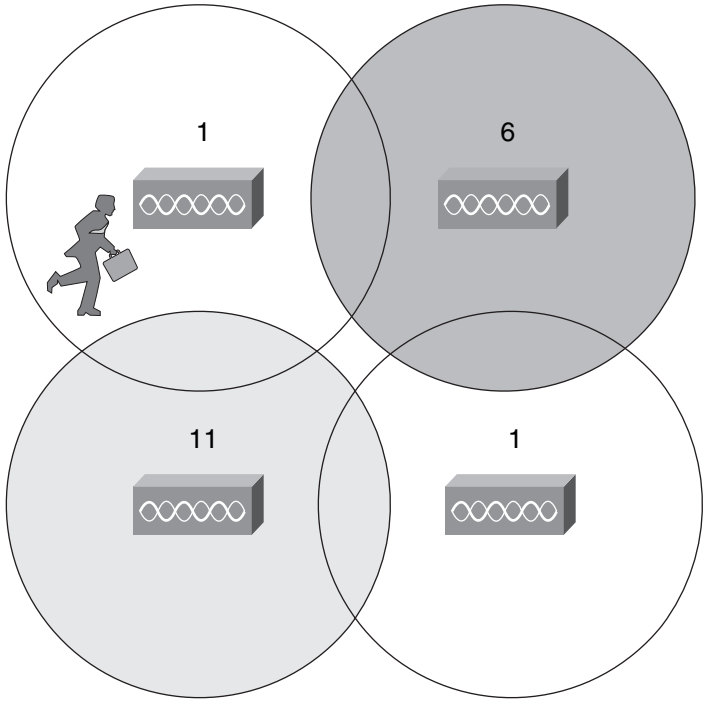

However, notice what is happening in the center where the cells meet—there is a small hole in RF coverage. If a client roams through that hole, his wireless signal will probably drop completely. As well, if the cells were brought closer together to close this hole, the two cells using channel 1 would overlap and begin interfering with each other.

Instead, you should lay out the cells in a "honeycomb" fashion as illustrated in Figure 18-7. This pattern is seamless, leaving no holes in coverage. In addition, notice how the two cells using

channel 1 are well separated, providing isolation from interference. As far as ordering channels in the pattern, several different variations are available using combinations of the three channels, but the result is basically the same.

Figure 18-7 *An Alternating Channel Pattern in 802.11b/g*

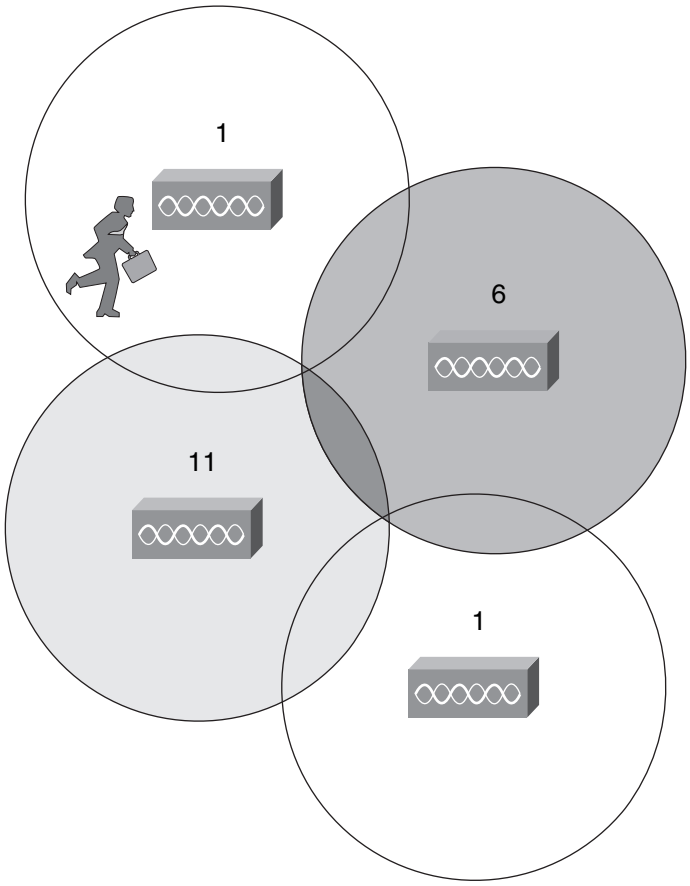

Notice that as the client shown in the channel 1 cell moves around, it will roam into adjacent cells on different channels. In order for roaming to work properly, a client must be able to move from one channel into a completely different channel.

Alternating channels to avoid overlap is commonly called *channel reuse*. The basic pattern shown in Figure 18-7 can be repeated to expand over a larger area, as Figure 18-8 illustrates.

Figure 18-8 *802.11b/g Channel Reuse over a Large Area*

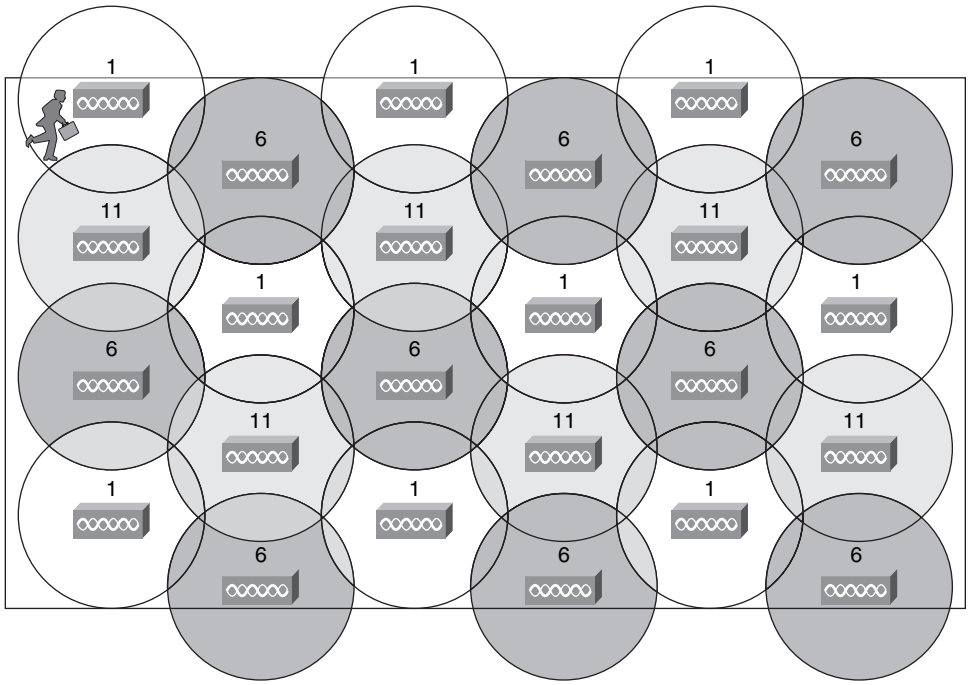

With 802.11a, the design is quite different. It has 4, 8, or even 12 nonoverlapping channels available, so the likelihood of adjacent cells using the same channel is very low. The FCC has added 11 additional channels in the U.S., for a total of 23 nonoverlapping choices.

So far, only the channel layout of a two-dimensional area has been discussed. For example, Figure 18-8 might represent only one floor of a building. What happens when you need to design a wireless LAN for multiple floors in the same building?

Recall that an RF signal propagating from an antenna actually takes on a three-dimensional shape. With an omnidirectional antenna, the pattern is somewhat like a donut shape with the antenna at the center. The signal extends outward, giving the cell a circular shape along the floor. The signal also extends upward and downward to a lesser extent—affecting AP cells on adjacent floors as well.

Consider the building with three floors shown in Figure 18-9. The same two-dimensional channel layout from Figure 18-8 is being used on the first floor. The floors in the figure are shown greatly separated, so that you can see the channel patterns and numbers. In reality, the cells on adjacent floors would touch or overlap, just as adjacent cells on the same floor do.

Figure 18-9 *Channel Layout in Three Dimensions*

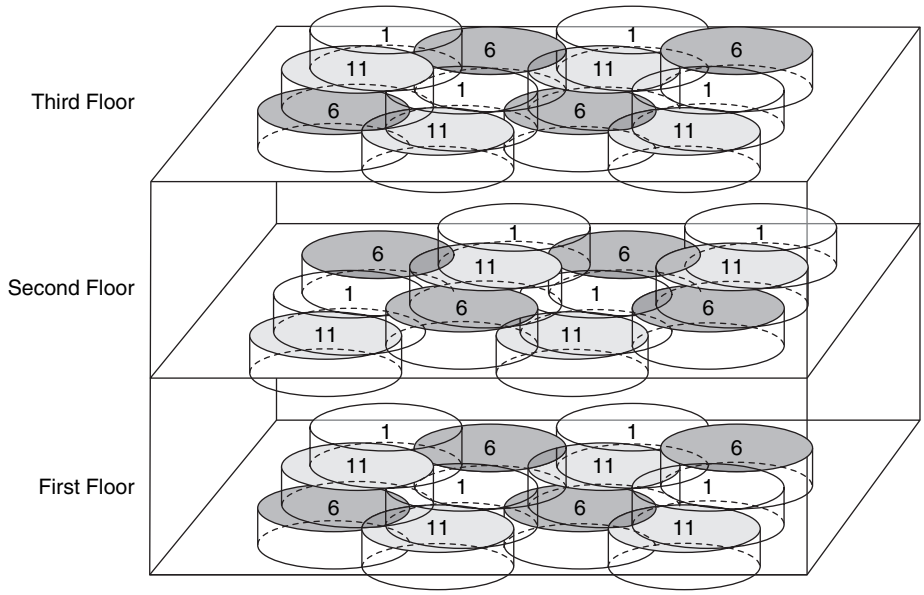

Now comes the puzzle of alternating channels within the plane of a floor, as well as between floors. Channel 1 on the first floor should not overlap with channel 1 directly above it on the second floor or below it in the basement.

When you consider each of the tasks involved in designing and maintaining a wireless LAN, it can really become a puzzle to solve! The cell size, AP transmit power, and channel assignment all have to be coordinated on each and every AP. Roaming also becomes an issue on a large scale, if clients are permitted to roam across the entire campus wireless network.

The good news is that Chapter 19, "Cisco Unified Wireless Network," explains how to solve many of these puzzles.

Foundation Summary

The Foundation Summary is a collection of information that provides a convenient review of many key concepts in this chapter. If you are already comfortable with the topics in this chapter, this summary can help you recall a few details. If you just read this chapter, this review should help solidify some key facts. If you are doing your final preparation before the exam, this information will be a convenient way to review the day before the exam.

Table 18-3 *Comparison of Wireless LAN Security Methods*

Method	Credentials Used	Data Security
Open Authentication	None	None
Pre-Shared Key (static WEP)	Static WEP key	WEP encryption (static key)
LEAP	Username/password	WEP encryption (dynamic keys)
EAP-TLS	Digital certificate	WEP encryption (dynamic keys)
PEAP	Server: Digital certificate Client: Any EAP method	WEP encryption (dynamic keys)
EAP-FAST	PAC to build tunnel; Username/password	WEP encryption (dynamic keys)
WPA	Any EAP-based method or pre-shared key	WEP encryption, per-packet keys generated by TKIP MIC performs packet authentication for integrity
WPA2	Any EAP-based method Pre-shared key Proactive key caching (authenticate once while roaming across APs)	AES encryption or TKIP MIC performs packet authentication for integrity

In this table, the darker shading indicates increasing levels of security.

Table 18-4 *Client Roaming Methods*

Roaming Method	Description
Passive roaming	Client scans other channels to listen for beacons from candidate APs
Active roaming	Client scans other channels and sends 802.11 probe requests to find candidate APs

Table 18-5 *Best Practice AP Cell Size Guidelines*

AP Radio	Typical Maximum Throughput Per Cell
802.11b	6.8 Mbps
802.11g	32 Mbps
802.11a	32 Mbps

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answers. Rather than allowing you to derive the answers from clues hidden inside the questions themselves, the questions challenge your understanding and recall of the subject. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

You can find the answers to these questions in Appendix A.

1. Consider the following wireless security methods. Based on encryption alone, which could be considered the most secure?
   a. Open authentication
   b. Pre-shared key
   c. LEAP
   d. WPA
   e. WPA2
2. Describe the weakness of using static WEP keys.
3. When is it considered appropriate to secure a WLAN with only open authentication, assuming that the SSID is kept secret?
4. What things should you consider when sizing an AP cell?
5. Should adjacent AP cells be configured to operate on the same channel? Why or why not?
6. Suppose that a set of 802.11b APs are configured to use channels 1, 6, and 11 in an alternating fashion. What channels should you choose when the 802.11g portion of the same APs are configured?
7. Suppose that 802.11a clients and APs are introduced into an area already using 802.11b or 802.11g. What effect will this have?
8. Suppose that an office area has wireless coverage from several APs, each using a very small cell size. Should you be concerned about the power output (and therefore the cell size) of client devices?