Part I

Concepts

In This Part

CHAPTER

1

Network Architecture Concepts

To fully appreciate the operation of IEEE 802.1X–based solutions, you must have a good understanding of network architecture concepts. It's important to know the component hardware that comprises a network, the layering process used by networking protocols, various IEEE 802 standards, and wireless network issues. This chapter provides an introduction to these basic concepts, which makes the process of learning 802.1X in later chapters much easier. Unless you're a seasoned network professional, you should read this chapter. If it's been awhile since you've covered networking concepts, then spend some time skimming through the chapter, and refresh yourself with the details that you might have forgotten.

Computer Network Defined

A *computer network* enables communications between computer-based client devices, servers, and peripherals to support various applications. Figure 1-1 illustrates a computer network. In terms of the client devices, the network appears as a cloud that provides the interconnection among client devices and servers. In realty, the cloud consists of a network infrastructure that provides routing of data, voice, and video from one point to another.

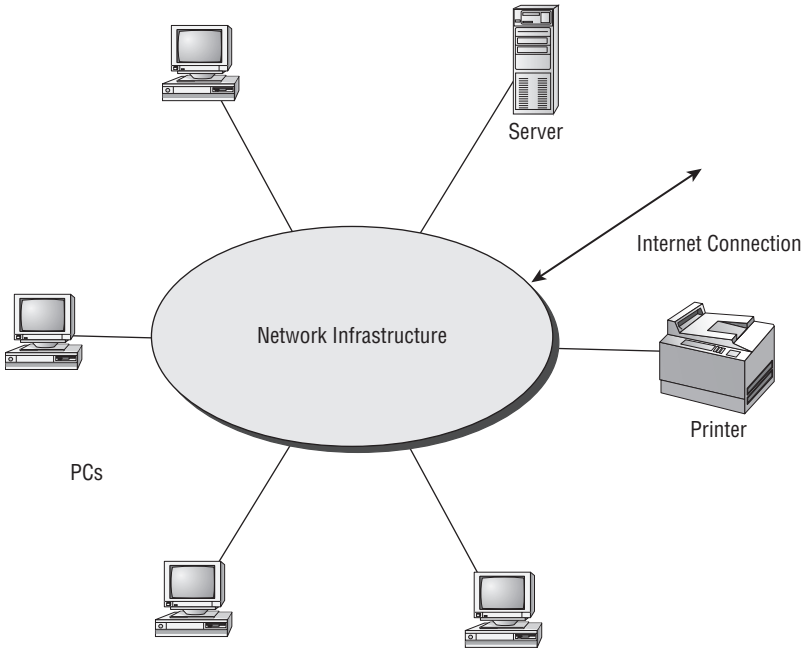

Figure 1-1: Computer network interconnecting client devices and servers

The most common network applications include e-mail, website browsing, file transfer, corporate applications, Internet Protocol (IP) telephony, and video-based security. We're all very familiar with e-mail and browsing the Web. With the use of File Transfer Protocol (FTP) software, you can also move files from one computer to another computer attached to the network. Of course many websites incorporate file-download capabilities for allowing users to download music, application software, and hardware firmware and software upgrades. Many enterprises also use networks to transport voice through the use of voice-over-IP (VoIP). In some cases, companies are making use of the network to interconnect security cameras and other video sources. As this book will explain in detail, IEEE 802.1X port-based security controls access to these applications, which can sometimes be mission-critical.

Network Components

Several basic components make up a computer network. These components include the following:

- Client devices
- Servers
- Network hardware

- Media
- Communications protocols

Client Devices

There are many different network client devices. Most users are familiar with using PCs, laptops, and printers on their network, but IP phones, cameras, game machines, and audio equipment are also becoming more common as client devices for networks. Figure 1-2 includes several examples of network client devices. IEEE 802.1X port-based authentication focuses on either allowing or not allowing client devices to have access to the network.

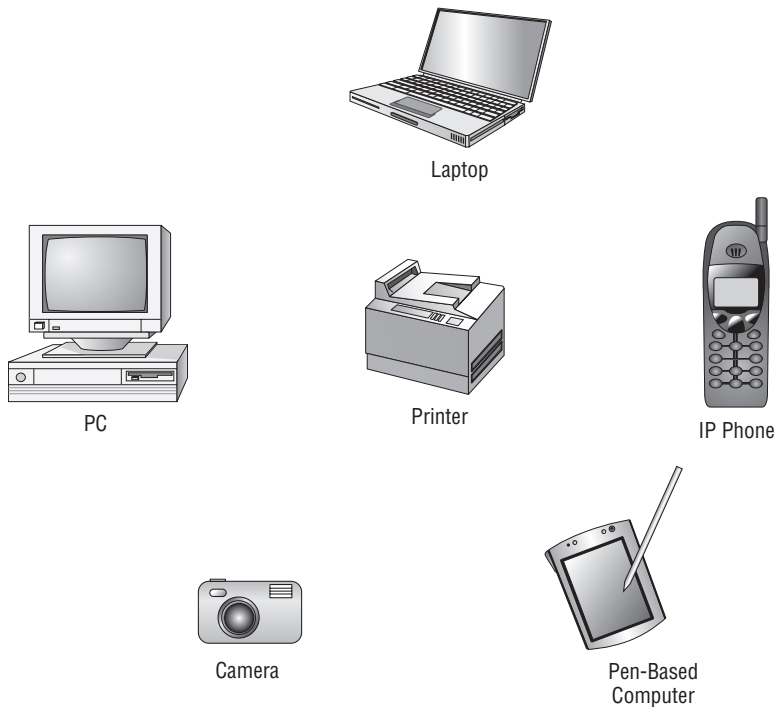

Figure 1-2: Various client devices

Servers

Network servers host application software and databases that users can access on the network. For example, an enterprise may have a warehouse management system that warehouse clerks and managers can access in order to perform inventories, check stock levels, and execute shipping of warehouse items.

In addition, as you'll see in Chapter 2, an authentication server will run *RADIUS* in order to process client devices that are authenticating with the network. A server is a centralized and shared component of the network (see Figure 1-3). As a result, sometimes the network may include multiple servers offering the same functionality in a redundant manner to increase availability by avoiding a single point of failure.

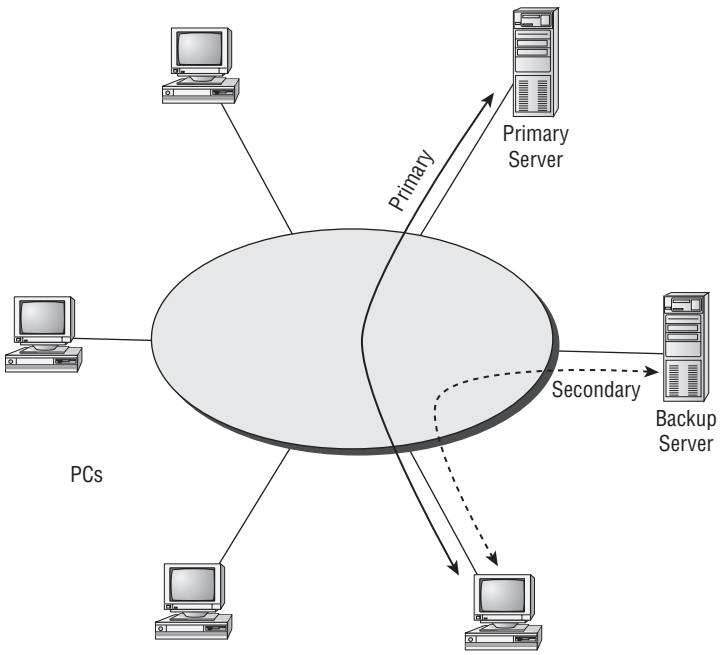

Figure 1-3: Backup servers provide higher availability.

The servers on the network consist of a hardware platform, which can be a PC running an operating system, such as Windows or Linux. The hardware platform includes a processor, memory, and interfaces that meet the performance requirements of the application that the server is supporting. The server hosts the actual application software, which an administrator installs and manages.

NOTE In addition to implementing IEEE 802.1X port-based authentication, administrators should be certain to lock down all administrative ports on servers and deactivate any un-needed utilities that the installation of application software spawns. This prevents a hacker, who is accessing the network servers from an authorized account, from manipulating the system and exploiting the network resources.

Network Hardware

The network hardware comprises the physical network infrastructure. All client devices and servers connect to the network infrastructure. Network hardware includes the following:

- Switches and hubs
- Routers
- Gateways
- Access points
- Network interface cards

Switches and Hubs

Switches and hubs provide wired interconnection points throughout the network infrastructure for the client devices. A switch or hub has multiple ports and implements the IEEE 802.3 standard (Ethernet) and IEEE 802.1 bridge protocols. Each port has a connector for attaching an Ethernet cable that connects to a client device (or other network hardware).

Figure 1-4(a) illustrates the function of a switch, which is slightly different than a hub. A switch, which implements what's referred to as *switched Ethernet*, connects a device on one port directly with another device on a different port. This connection doesn't preclude devices connected to other ports from communicating with the switch and forming connections with other devices. After the connection between two devices is made inside the switch, other devices can contend for connections.

A switch offers much faster performance than a hub, as Figure 1-4(b) depicts. When a device connected to a port on a hub connects a device connected to a different port on the same hub, the devices connected to the other ports on the hub can't communicate with the hub and form other connections. All devices interfacing with the ports on a hub, except the two already connected with each other, are blocked. As a result, most enterprises use switches.

In general, switches are relatively simple devices. They forward data, including multicast packets, throughout the network based on the Media Access Control (MAC) address that each packet is carrying. The MAC address is the actual physical address that the network devices respond to, similar to the number and street address of your home or office.

NOTE Some servers and applications periodically transmit broadcast packets, and switches forward these packets through, which can sometimes cause unnecessary traffic on the network.

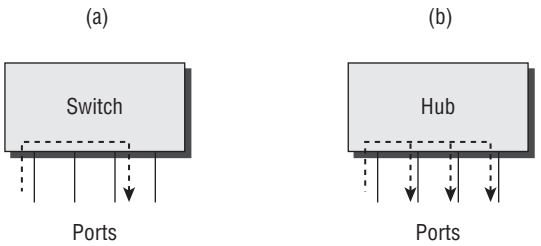

Figure 1-4: Ethernet switch vs. a hub

Routers

A router is a bit more sophisticated than a switch or hub. Routers actually route packets through a network infrastructure, as Figure 1-5 illustrates, based on the IP address carried within each packet. The IP address corresponds to the addressing that many applications respond to. As packets flow into one particular port, the router will make a decision on which outbound port to use—that is, the one that will move the packet closest to the intended destination. A *routing protocol* keeps track of which outbound ports provide optimum routing and maintains this information in a routing table. The router simply looks at the destination IP addresses the packet is carrying and looks up this destination in the routing table. The routing table contains an outbound port entry for each possible destination.

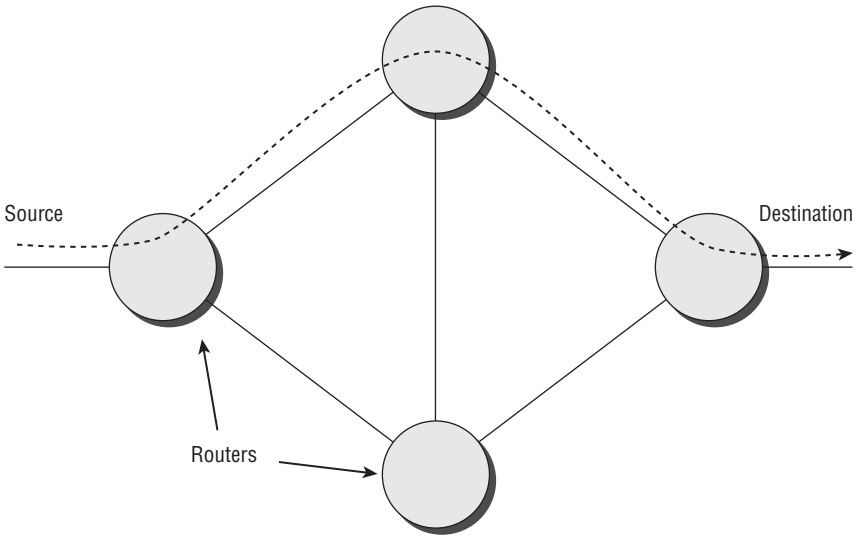

Figure 1-5: Routers route packets through the network.

An enterprise network infrastructure will generally consist of switches at the edge of the network, interfacing client devices to the network, and routers that interconnect groups of switches into separate subnets. This organization increases performance and eases management of the system.

Access Points

Wireless networks use access points to create a radio cell where wireless client devices can associate with the access point, which enables the wireless client devices to send and receive data over the wireless network. An enterprise or city will often have dozens, hundreds, or even thousands of access points installed at locations so that the radio cells of adjacent access points overlap slightly. The multiple overlapping radio cells create continuous signal coverage so that wireless client devices can roam from an area covered by one access point to an area covered by a different access point.

With access points, the flow of traffic between one wireless client and another wireless client associated with the same access point must flow through the access point, as shown in Figure 1-6. Traffic flowing from a wireless client and another network device connected to the network infrastructure flows through the access point as well. As with switches and hubs, access points implement the IEEE 802.1 bridge protocols.

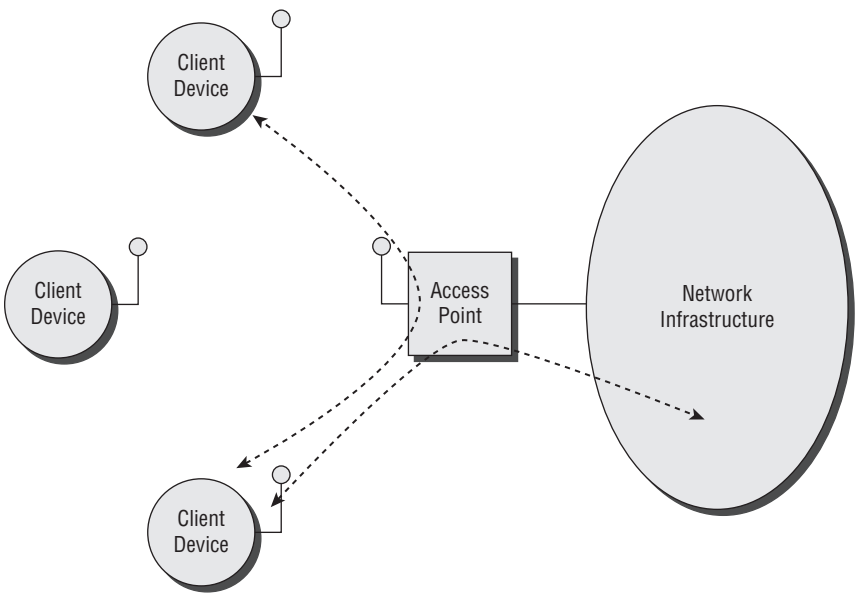

Figure 1-6: Access points interface wireless client devices to the network.

Access points are similar to hubs, except that the connections made with the client devices are wireless. Thus, an access point translates between wireless network protocols and wired network protocols. With IEEE 802.11 access points, multiple client devices can associate with a single access point at the same time, but, similar to the hub, only a single client device (or the access point) can transmit at any given time.

A switch or hub interconnects multiple access points to form a wireless network infrastructure. Each access point connects to a different wired port on the switch or hub. In most cases, the port provides Power over Ethernet (PoE) for supplying electrical power to the access point. Electrical wiring does not need to be installed at each access point location—all that's necessary is an Ethernet cable that connects the access points, limited to approximately 300 feet.

Some access points also include routing functions, and are commonly referred to as wireless routers. A wireless router is suited for home or small office applications where it's beneficial to have a wireless local area network (LAN) provided by a single device that connects to the Internet via a digital subscriber line (DSL) or cable modem. The wireless router performs the functions of an access point and also includes Dynamic Host Configuration Protocol (DHCP) and Network Address Translation (NAT). This enables the wireless network to use a single official IP address provided by the Internet service provider (ISP) for multiple client devices.

NOTE To maximize interoperability throughout an 802.1X system, be sure to keep the firmware up-to-date on all network hardware.

Network Interface Cards

Network interface cards (NICs), also referred to as client cards or network adapters, interface the client device (and other network devices) to the network. For example, you can interface a laptop to a wireless network by installing an 802.11 PC Card in the laptop. The NIC is the actual network connection, addressable by a MAC address. The following sections describe each of the various NIC form factors, which defines the physical interface between the network card and the computer.

ISA

The Industry Standard Architecture (ISA) bus is the most common bus interface in the desktop PC world. ISA has been around since the early '80s for use in the IBM PC/XT and PC/AT. Because of this, the proliferation of the ISA has been significant. Despite its lack of speed (2 Mbps), nearly all PCs manufactured up until several years ago had at least one ISA bus. The ISA bus has

failed, however, to advance at the pace of the rest of the computer world, and other higher speed alternatives are now available. ISA doesn't impose too much of a performance impact on IEEE 802.11b wireless LANs, but it's not advisable to purchase new ISA cards because they may become obsolete.

PCI

The Peripheral Component Interconnect (PCI) bus is the most popular bus interface for PCs today and boasts a throughput rate of 264 Mbps. Intel originally developed and released PCI in 1993, and it satisfies the needs of most recent generations of PCs for multimedia, graphics, and networking cards.

PCI cards were the first to popularize Plug and Play (PnP) technology. PCI circuitry recognizes compatible PCI cards and then works with the computer's operating system to set the resource allocations for each card. This helps save time and prevents installation headaches.

PC Card

Developed in the early '90s by the Personal Computer Memory Card International Association (PCMCIA), the PC Card is a peripheral device (about the size of a credit card) that can provide extended memory, modems, connectivity to external devices, and, of course, wireless LAN capabilities to laptops. PC Cards are the most widely available cards for portable devices, such as laptops. In fact, they are more popular than ISA or PCI cards because of a growing usage of laptops.

If you want to share a PC Card with your desktop PC, consider using an adaptor that converts a PC Card into a PCI card. This allows you to purchase one card for use in both types of computers. You can take the PC Card on your business trips or home from work and use the same card when you're back in your office on a PC.

Mini-PCI

A Mini-PCI card is a small version of a standard desktop PCI card. It has all the same features and functionality of a normal PCI card, but it's about one quarter the size. Mini-PCI cards are integrated within laptops as an option to buyers, with antennas that are often integrated within the monitor's case or even next to the LCD screen. A strong advantage of this form of radio NIC is that it frees up the PC Card slot for other devices. In addition, manufacturers can provide Mini-PCI–based wireless connectivity at lower costs.

CF

SanDisk Corporation first introduced CompactFlash (CF) in 1994, but wireless LAN radio cards were not available in CF form until recently. A CF card is very small—it weighs half an ounce and is less than half the thickness and one

quarter the volume of a PC Card radio card. The CF cards also draw very little power, which enables batteries to last longer than on devices that use PC Cards. Some PDAs (Personal Digital Assistants) come with direct CF interfaces, which results in a very lightweight and compact wireless PDA. A CF radio card is definitely the way to go, especially for compact computing devices.

NOTE Consider using a CF-to-PC Card adapter to operate the CF card in a laptop via a PC Card slot. This enables you to take advantage of the miniature radio card in the PDA and not have to purchase another card for your laptop.

Media

The media in a network interconnects all of the network components. Networks use the following types of media:

- Wire
- Optical fiber
- Air

Metallic Wire

The use of metallic wires is the most common method for interconnecting components. When computer networks first came into existence, most network wiring was coaxial cable, which was bulky and difficult to install. Conventional network wiring today consists of *twisted pair cabling,* with the most common being Category 5 twisted pair. Digital data, video, and voice signals in the form of electrical current run through the wires. Twisted pair cabling can support data rates into the Gigabit per second (Gbps) range. IEEE 802.3 (Ethernet) specifies the use of twisted pair wiring.

The flow of electrical current through metallic wires, such as Category 5 twisted pair cabling, creates an electromagnetic field around the wire. With sensitive listening equipment, a hacker could sense this field from several feet away from the wiring and possibly decode the applicable data. However, this would generally require the hacker to gain access to the facility—eavesdropping on a wired network from outside the building would be difficult, if not impossible.

Nearly all businesses have Category 5 cabling running throughout facilities, with end points (network taps) located in most offices. The cables run from each office to one of several wiring closets located in a facility. The wiring closets house the switches. The range of Category 5 cabling is limited to less than approximately 300 feet. The wiring closets are where the company places switches and

other network equipment. Generally, a higher-speed backbone interconnects the switches located in the wiring closets as shown in Figure 1-7.

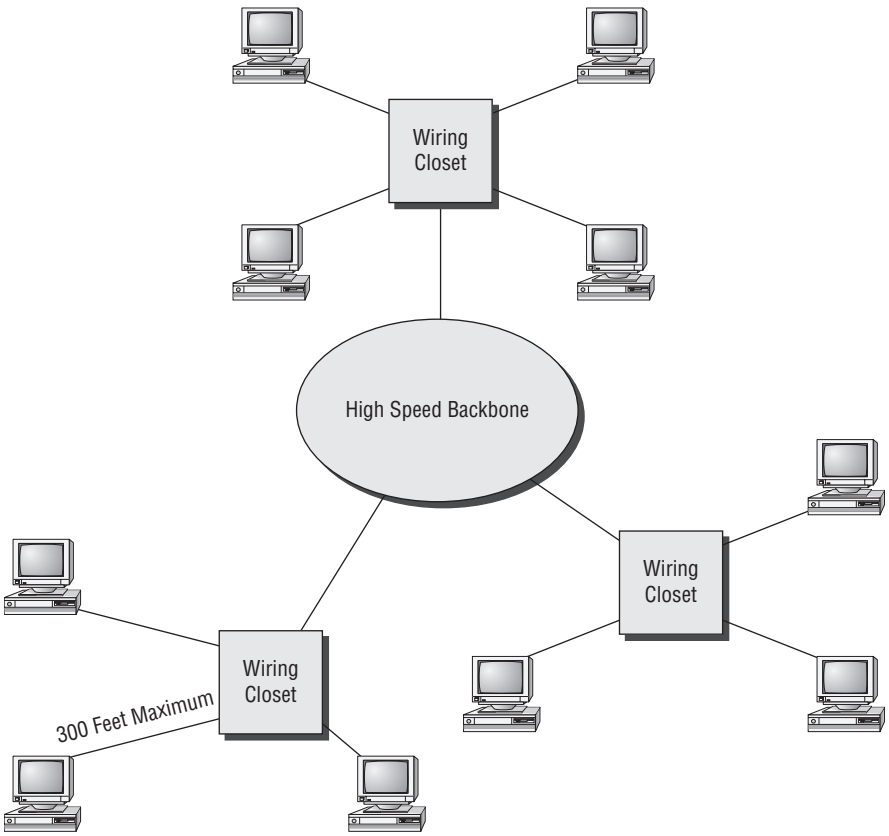

Figure 1-7: Typical wiring in an enterprise wired network

Optical Fiber

Optical fiber cabling consists of strands of glass that conduct light efficiently from one end to the other. The advantage of optical fiber as compared to metallic wiring is that fiber supports much higher data rates (up to tera bits per second), doesn't emit electromagnetic fields, and operates over a relatively long range. Thus, optical fiber makes the system more secure and supports higher-speed delivery of information. In addition, optical fiber is practical for providing high-speed connections between buildings throughout cities. An issue with optical fiber, though, is that is relatively expensive to install for each client device inside a company. In some cases, however, especially where client devices are beyond the 300-feet range from a wiring closet, optical fiber may be the only way to connect the clients.

Air

Air is the medium for wireless systems, which make use of radio waves or infrared light for transporting data, video, and voice signals. Radio waves, generally in the 2.4-GHz or higher frequencies, are the most common signaling for wireless networks. Most of the IEEE 802 standards (such as 802.11a, 802.11b, 802.11g, 802.11n, 802.15, and 802.16) use radio waves. There is an infrared version of 802.11, but there are very few, if any, implementations.

An issue with using air as a medium is that obstacles get in the way. A building will certainly have walls between wireless client devices and access points. In outdoor environments, trees and buildings fall in the path of radio and light signals. Radio waves can travel through most materials, with varying attenuation. Infrared light doesn't penetrate most obstacles it encounters. Thus, when installing wireless systems, you must carefully analyze the environment, and design the system to accommodate attenuation that occurs between the client devices and the access points. This can be difficult to accomplish, because obstacles may change position.

Electromagnetic signal sources, such as microwave ovens, other wireless systems, and even sunspot activity can impact the performance of a wireless network. These sources of interference can be sporadic and unpredictable, which may affect the network's ability to meet specified service levels. As a result, client-device connectivity may be interrupted from time to time and disrupt the data flow and protocol operation between the clients and the network.

Infrared light can't be seen by humans in most lighting conditions, but it's free from electromagnetic interference and can operate at relatively high data rates over a range of one to two miles. Cities will often implement a point-to-point infrared system to interconnect buildings, rather than undertake an expensive cable installation.

Network Types

There are several different types of computer networks, categorized primarily by the size of the area that the network operates. The following sections explain each type of network.

Personal Area Networks

A *personal area network* (PAN) connects computer devices within a relatively small area, such as in the immediate vicinity of a person's body. Figure 1-8 illustrates a PAN. A PAN may require wires for connecting the devices, or connectivity may be wireless. For example, a wired PAN may consist of a smart

phone connected to a laptop via a USB cable, and headphones wired to the audio output jack on the laptop. A wireless PAN, however, may connect a phone, laptop, and headphones wirelessly. Both wired and wireless PANs allow users to listen to music streaming from a laptop or use a contact manager on the laptop to place telephone calls. Or, a person may use a PAN to connect their smart phone to a PC to synchronize contacts, schedules, and tasks.

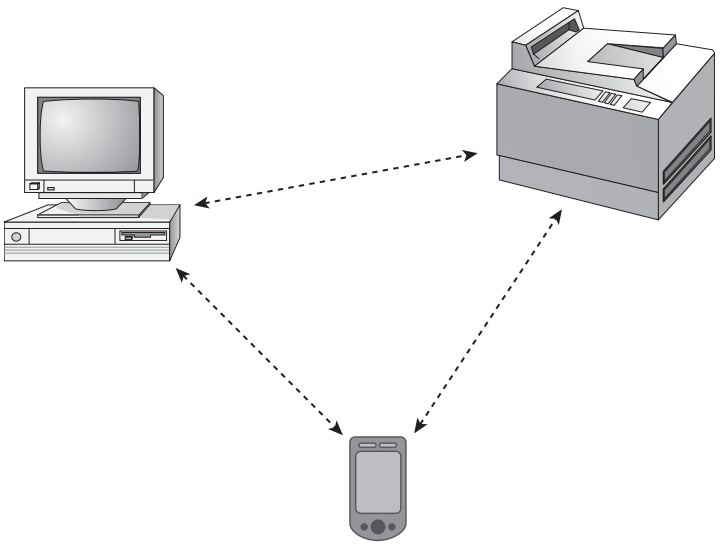

Figure 1-8: PANs connect computer devices over limited areas.

REDUCING CABLES AND IMPROVING CONVENIENCE WITH A WIRELESS PAN

Ralph purchases a new smart phone that includes a Bluetooth radio. After synchronizing the phone with a PC in his office a few times, he finds that the cable is too cumbersome. In order to make the synchronization process go smoother, he creates a wireless PAN.

Ralph purchases a Bluetooth dongle that attaches to the USB port on the PC. This allows the smart phone to synchronize with the PC without connecting a cable. In fact, he sets the phone to automatically synchronize every few minutes. With this configuration, Ralph doesn't have to remember to attach a cable and synchronize his phone before leaving his office. The phone keeps up-to-date wirelessly while sitting on the desk or holster. After Ralph gets this working, he discovers that he can add a wireless ear pod to his wireless PAN in order to listen to music streaming from his PC or phone. He uses this setup for hands-free operation of his phone while driving a car.

Most PAN networking technologies, such as IEEE 802.15 (Bluetooth wireless PAN) and other proprietary technologies, provide point-to-point network connectivity. In fact, many of the Bluetooth wireless dongles attach to a USB port on a PC or laptop and interface two devices, such as a PC and smart phone, via a serial port. In this manner, the wireless dongle is merely replacing the serial cable with a wireless connection.

Local Area Networks

A *local area network* (LAN) connects computer devices that span the size of a building or college campus. An enterprise facility or hospital will likely have a LAN that connects PCs and corporate servers. A warehouse clerk, for example, may use a PC connected to a LAN to find the status of received goods or set up a shipping transaction. In addition, a LAN will generally connect to the Internet, allowing users of the LAN to browse websites and send e-mail to users who are not connected to the LAN.

Figure 1-7 depicts a wired LAN. Nearly all companies and organizations have wired LANs, which are generally based on IEEE 802.3 (Ethernet) technologies. With Ethernet, users' computer devices connect to the network via an Ethernet cable. Ethernet is very common because it's been in existence for over 20 years. You'll find wired Ethernet LANs in nearly all facilities, such as enterprises, hospitals, and universities.

One issue with wired LANs is that they require users to operate their computers from stationary locations. This isn't much of a problem with most PCs, mainly because of their size. You simply plug a PC in to the Ethernet network and don't move it unless you're moving from one office to another. PCs don't usually need the benefits of mobility that wireless connectivity offers, except for some unique and relatively uncommon applications. For example, in an elementary school that has limited funding and can only purchase one PC for multiple classrooms, the PC can be placed on a rolling cart and be wirelessly interfaced to the network—enabling teachers to easily move the PC from classroom to classroom for effective sharing.

IEEE 802.11 wireless LANs are beginning to become more commonplace in businesses and homes. For facilities that already have wired LANs, adding wireless connectivity through a wireless LAN extends networking to areas not covered by the wired LAN (such as warehouses) and offers mobility for users (see Figure 1-9). In addition, some facility owners and managers elect to install a wireless LAN instead of wired Ethernet as the primary network for the entire facility. This isn't too common, however.

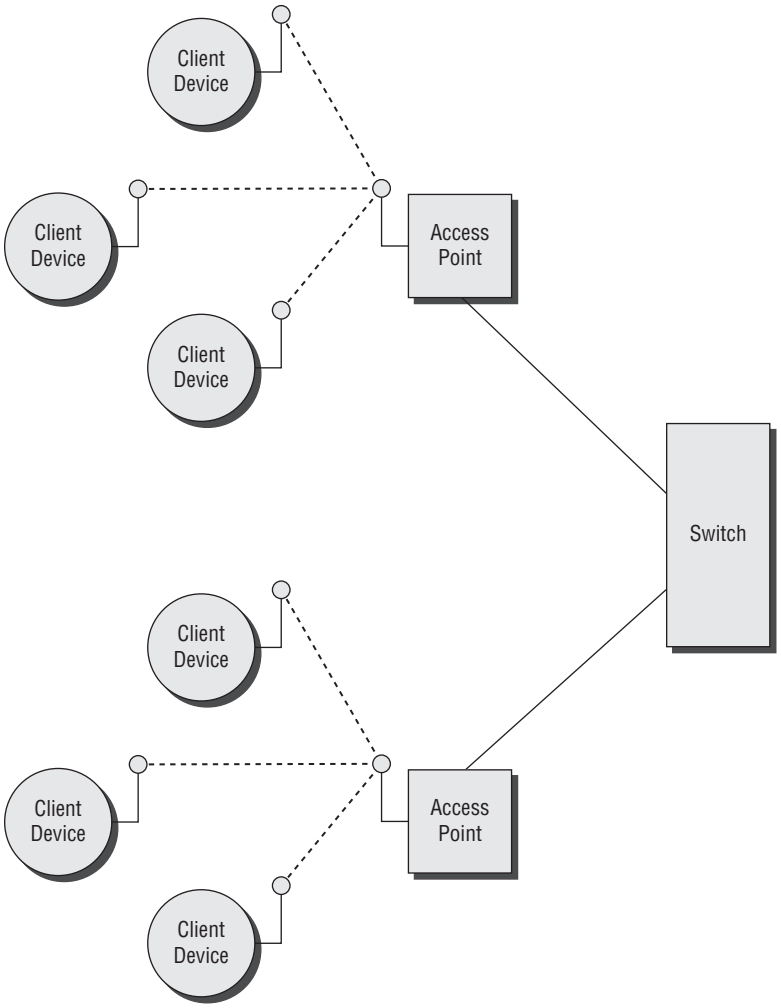

Figure 1-9: A wireless LAN enables mobility.

NOTE **Wi-Fi, which stands for *Wireless Fidelity*, was created by the Wi-Fi Alliance based on IEEE 802.11 standards. The Wi-Fi Alliance requires that products undergo interoperability testing with other vendor products in order to bear the Wi-Fi logo.**

ADDING WIRELESS CONNECTIVITY TO AN EXISTING WIRED LAN TO SUPPORT WIRELESS IP PHONES

A major hospital in the Midwestern United States already has a wired Ethernet network with a LAN that supports hundreds of PCs throughout the hospital. The existing system does a great job of supporting e-mail and remote access to applications, but doctors, nurses, and administrative staff can benefit by having mobile phones. Dependence on wired phones often causes significant delays in coordinating patient care as a result of the "telephone tag" that happens when staff are only sporadically near a phone.

As a result, the hospital deployed a wireless LAN throughout its facilities to support wireless IP phones. This was a much less costly alternative to paying for cellular service within the hospital. With mobile phones, the hospital is able to react much faster to patient emergencies. In addition, the deployment of the wireless LAN puts the hospital in a position to use other wireless applications, such as electronic patient records.

Metropolitan Area Networks

A *metropolitan area network* (MAN) provides network connectivity over the area of an entire city or metropolitan area. For cities such as Houston and Los Angeles, a MAN can span hundreds of square miles. In other places, such as smaller cities, a MAN may cover only a few square miles.

Optical Fiber Infrastructure

Many local governments have optical fiber installed under streets throughout most of their downtown areas. This fiber infrastructure is in place to provide reliable high-speed connections between various buildings throughout the city. Fiber taps are available in many locations to facilitate interfacing network equipment that may operate at a particular location. For example, some cities connect wireless access points or mesh nodes to the fiber taps and make use of the existing fiber infrastructure to provide network connectivity. Or, the fiber infrastructure may carry network traffic from inside a building, such as city hall, to a centralized data center for connection to the Internet.

Wi-Fi Mesh

Many municipalities are installing *Wi-Fi mesh* networks to provide wireless connectivity for mobile Wi-Fi–equipped client devices. For example, a city may deploy a mesh network to support wireless public-safety applications or wireless work-order processing. In addition, a municipal Wi-Fi network can provide public Internet access throughout the city.

Figure 1-10 illustrates the physical architecture of a Wi-Fi mesh network. A mesh node is similar to a wireless access point—it offers a Wi-Fi interface to client devices based on the IEEE 802.11 protocol. However, unlike access points, mesh nodes don't require a cable that connects the node to a switch. Instead, each mesh node is capable of receiving packets from a wireless client device and forwarding the packets to another mesh node. The mesh nodes implement a routing mechanism that moves packets toward a mesh node with a link back to a central connection point. Some of the mesh nodes are also gateways, which generally use a point-to-point radio link (called *backhaul*) that connects the gateway to the central point. For client devices connecting to a mesh node (not a gateway), a data packet hops from one mesh node to another mesh node until the packet arrives at a gateway, which then sends the packet directly to the central point.

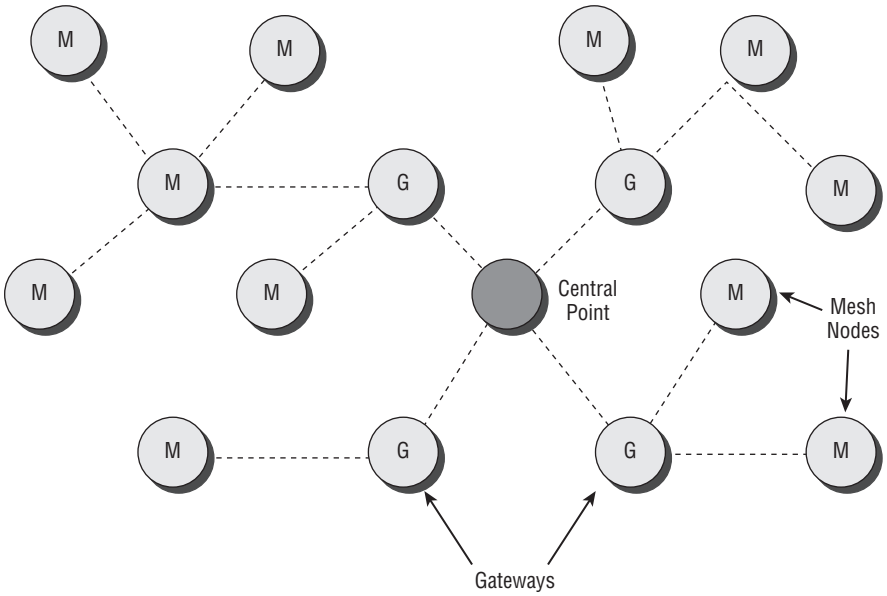

Figure 1-10: Mesh nodes allow packets to hop through the network without wires.

WiMAX

WiMAX (which stands for *Worldwide Interoperability for Microwave Access*) is a relatively new specification issued by the WiMAX Forum for providing wireless network connectivity within metropolitan areas. WiMAX is based on the IEEE 802.16 standard. The most current version of the standard, IEEE 802.16-2005, addresses both fixed and mobile wireless connectivity. Many of the WiMAX

deployments to-date use licensed frequency spectrum. For example, Sprint Nextel is in the process of deploying mobile WiMAX in many major U.S. cities. Some cities that deploy Wi-Fi mesh networks make use of fixed WiMAX for connecting mesh nodes (gateways) to central points in the system. As more and more client devices become equipped with WiMAX, it may eventually replace Wi-Fi as the technology of choice for wireless MANs.

Wide Area Networks

Wide area networks (WANs) cover very large areas, sometime the entire globe. Companies with offices spread throughout a nation or the entire world can make use of a WAN to interconnect the individual networks in offices. The Internet is actually a WAN that many companies depend on every day for supporting e-mail, file transfer, and access to remote applications. In some cases, a company may create a private WAN. For example, a retail company may link all of its stores to a central data center. Every night, each of the stores can upload sales data and download changes, such as price updates. WANs make use of long-haul leased terrestrial or satellite links.

Logical Network Architecture

As discussed earlier in this chapter, a computer network consists of various hardware components, including client devices such as PCs, laptops, servers, switches, routers, and media. Up until now, this chapter addressed the network as a physical entity. In order to fully understand how the protocols work and interoperate, it's advantageous to examine a network based on its logical architecture. This view of the network is based on the classic seven-layer *Open Systems Interconnection (OSI) reference model* shown in Figure 1-11.

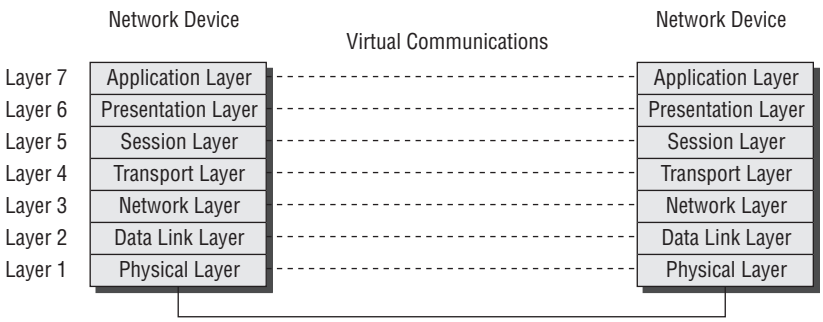

Figure 1-11: The OSI reference model offers seven layers of functionality.

Each layer of the model represents specific functions. A network may not implement all of the possible functions, but the model covers most of them. Each layer of the model does work for the layer above and doesn't care what happened at lower layers. In fact, each layer has no idea that lower layers are doing anything. In some cases, standards (such as IEEE 802.3 and IEEE 802.11) will subdivide layers of the OSI model into multiple sublayers as well.

The following explains the purpose of each layer:

- **Layer 7, Application Layer**: The application layer provides basic user applications, such as e-mail and file transfer. Simple Mail Transfer Protocol (SMTP) and File Transfer Protocol (FTP) are examples of Application Layer protocols.
- **Layer 6, Presentation Layer**: The Presentation Layer provides syntax for network entities to communicate.
- **Layer 5, Session Layer**: The Session Layer manages the flow of data between network entities.
- **Layer 4, Transport Layer**: The Transport Layer establishes and maintains end-to-end connections between network entities. Transmission Control Protocol (TCP) is an example of a Transport Layer protocol.
- **Layer 3, Network Layer**: The Network Layer provides routing throughout the network for data packets going from one network entity to another. The Internet Protocol (IP) is an example of a Network Layer protocol.
- **Layer 2, Data Link Layer**: The Data Link Layer establishes and maintains link-level connections between network devices, such as the link between an Ethernet card and an Ethernet switch. The IEEE 802 standards (such as 802.3 and 802.11) implement Data Link Layer functions.
- **Layer 1, Physical Layer**: The Physical Layer defines the electrical and mechanical specifications of the interface between network devices. For example, IEEE 802.11g is a Physical Layer standard that specifies the use of radio waves at 2.4 GHz.

NOTE The IEEE 802.1X and related standards apply to only Layer 2, the Data Link Layer.

Each specific OSI layer provides communications, relevant to its specific functions, between network entities. However, this is virtual rather than direct communication—the lower layers encapsulate data from layers above and deliver the encapsulated data to the opposite network entity. Chapter 2 explains this encapsulation process as it applies to 802.1X.

IEEE 802 Standards

The IEEE 802 standards play a big role in 802.1X port-based authentication systems. 802.1X is certainly part of the 802 standards. In addition, 802.3 and 802.11 are very important because they carry the authentication traffic.

Figure 1-12 illustrates the logical network model for the 802 standards. In general, 802.1, which includes 802.1X, provides overall network architecture for LANs, network management, internetworking between 802 LANs, and overall 802 security. 802.2, Logical Link Control (LLC), provides functionality similar to Layer 2, the Data Link Layer, of the OSI reference model. The Medium Access Control (MAC) enables multiple client devices to share a common medium, such as twisted pair cabling or a radio signal through the air. The 802.3 (Ethernet) and 802.11 (Wi-Fi) standards both define a MAC Layer. The Physical Layer (PHY) performs the actual transmission and reception of data. 802.3 100Base-T and 802.11g are PHY standards.

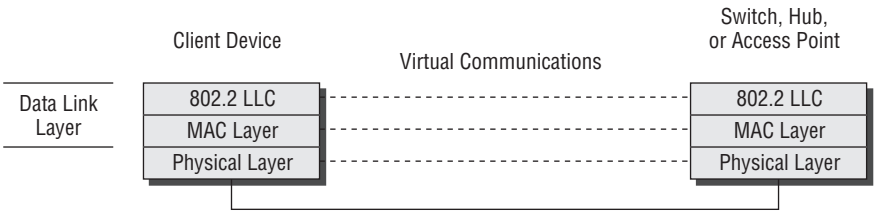

Figure 1-12: IEEE 802 logical architecture

There are several 802.3 PHY standards, which include various implementations of twisted pair wiring (10Base-T, 100Base-T, 1000Base-T) and fiber cabling. The 802.11 PHY standards include the following:

- **802.11a**: 5-GHz frequency band with data rates up to 54 Mbps; not compatible with 802.11b or 802.11g.
- **802.11b**: 2.4-GHz frequency band with data rates up to 11Mbps.
- **802.11g**: 2.4-GHz frequency band with data rates up to 54Mbps; backward compatible with 802.11b.
- **802.11n**: 2.4-GHz frequency band with data rates up to 100 Mbps; backward compatible with 802.11b and 802.11g.

Wireless Impairments

You should be familiar with the following wireless impairments when deploying 802.1X port-based authentication systems:

- Roaming delays
- Coverage holes
- Radio frequency (RF) interference

Roaming Delays

Wireless networks provide mobility that requires client devices to roam from one access point to another access point as the user moves about a facility or city. The Wi-Fi NIC in a client device periodically scans the immediate area for access points. When certain conditions are met, such as when the currently associated access point has relatively low signal strength and there are excessive data frame retransmissions, the client device NIC will initiate roaming from the currently associated access point to an access point with stronger signal strength. The client device re-associates with the new access point, and all corresponding data traffic is redirected through the new access point. Figure 1-13 illustrates the roaming process.

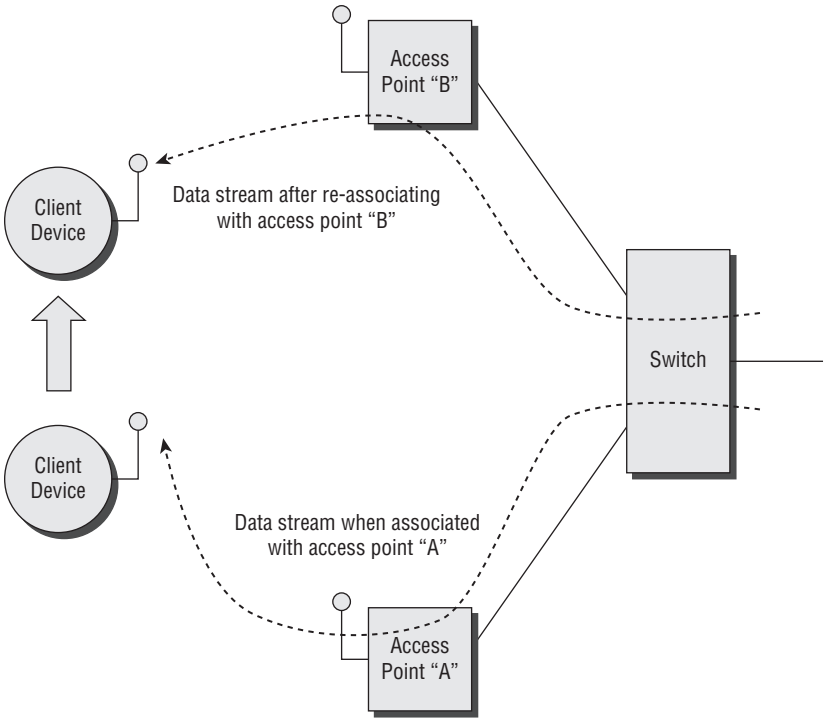

Figure 1-13: Wireless LAN roaming process

Ideally, this doesn't impede the flow of data between the client device and the network. The problem, however, is that some wireless NICs incorporate rather lengthy delays when roaming. This can wreak havoc on port-based authentication because some systems require the client device to re-authenticate at every newly associated access point. The roaming delays caused by the client device NIC may interrupt and even disrupt 802.1X authentication protocols, unless 802.1X configurations compensate for the delays.

REAL-WORLD ROAMING ON WIRELESS LANS

An extremely beneficial aspect of Wi-Fi networks is mobility. For example, a person can walk through a facility while carrying on a conversation over a Wi-Fi phone or when downloading a large file from a server. Ideally, the Wi-Fi radio inside the user device automatically roams from one access point to another as needed to provide seamless connectivity. But is this what really happens?

In the past, I've experienced issues with roaming, so I decided to perform some testing. I was especially curious about how fast roaming actually works and whether or not it's disruptive to wireless applications.

My test configuration included two access points: one access point (AP-1) set to channel 1, and the other access point (AP-2) set to channel 6. Other settings were default values, such as a beacon interval of 100 milliseconds, RTS/CTS disabled, and so on. The access points were installed in a typical office facility in a manner that provided a minimum of 25dB signal-to-noise ratio throughout each access point's radio cell, with about 20-percent overlap between cells. This is somewhat the industry standard for wireless voice applications. The roaming client in my case, though, was a laptop equipped with an internal Centrino Wi-Fi radio (Intel PRO/Wireless 2915ABG).

While standing with the wireless client within a few feet of AP-1, I used an AirMagnet Laptop Analyzer (via another Wi-Fi card inserted into the laptop's PCMCIA slot) to ensure that that I was associated with AP-1. I then kicked off an FTP transfer of a large file from the server to the laptop and started measuring the 802.11 packet trace using the AirMagnet Laptop Analyzer. With the file downloading throughout the entire test, I walked toward AP-2 until I was directly next to it. With the packet trace, I was able to view the exchange of 802.11 frames, calculate the roaming delay, and see if there was any significant disruption to the FTP stream.

After the client radio decided to re-associate, it issued several 802.11 disassociation frames to AP-1 to initiate the re-association process. The radio then broadcasted an 802.11 probe request to get responses from access points within range of the wireless client. This is likely done to ensure that the client radio has up-to-date information (beacon signal strength) of candidate access points prior to deciding which one to re-associate with.

AP-2 responded with an 802.11 probe response. Because the only response was from AP-2, the client radio card decided to associate with AP-2. As expected, the association process with AP-2 consisted of the exchange of 802.11 authentication and association frames (based on 802.11 open system authentication).

The re-association process took 68 milliseconds, which is the time between the client radio issuing the first dissociation frame to AP-1 and the client receiving the final association frame (response) from AP-2. This is quite good, and I've found similar values with other vendor access points.

The entire roaming process, however, will interrupt wireless applications for a much longer period of time. For example, based on my tests, the FTP process halts an average of five seconds prior to the radio card initiating the re-association process (issuing the first disassociation frame to AP-1). I measured 802.11 packet traces, indicating that the client radio card re-retransmits data frames many times to AP-1 (due to weak signal levels) before giving up and initiating the re-association with AP-2. This substantial number of retransmissions disrupted the file download process, which makes the practical roaming delay in my tests an average of five seconds. The Centrino radio card I tested is notorious for this problem, but I've found this to be the case with most other radio cards as well.

Vendors are likely having the radio cards hold off re-associations to avoid premature and excessive re-associations (access point hopping). Unfortunately, this disrupts some wireless applications. If you plan to deploy mobile wireless applications, then be sure to test how the roaming impacts the applications.

Every model radio card will behave differently when roaming due to proprietary mechanisms, and some cards will do better than others. Just keep in mind that roaming may take much longer than expected, so take this into account when deploying wireless LAN applications—especially wireless voice, which is not tolerant of roaming delays that exceed 100 milliseconds.

Coverage Holes

Changes made inside the facility after the initial wireless LAN is installed may alter radio frequency (RF) signal propagation. For example, a company may construct a wall, which offers significant attenuation that wasn't there before. Or, a thorough site survey may not have been done prior to installing the network. These situations often result in areas of the facility having limited or no RF signal coverage, which decreases the performance and disrupts the operation of wireless applications. In addition, the coverage holes, similar to roaming delays, may disrupt 802.1X port-based authentication system protocols, which can lead to some pesky issues. Figure 1-14 illustrates coverage holes.

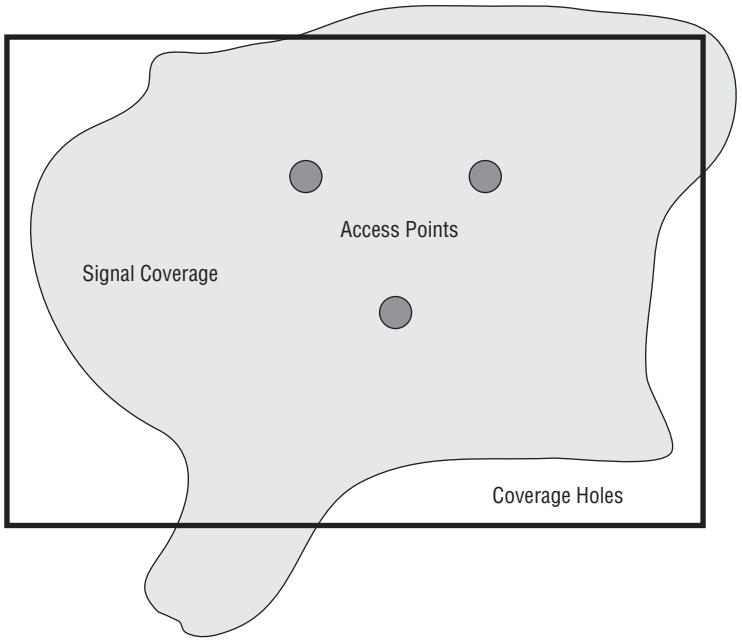

Figure 1-14: Wireless LAN coverage holes

Indications of a coverage hole include low signal-level (less than –75 dBm) and high retry rates (greater than 10 percent), regardless of noise levels. The signal in this situation is so low that the receiver in the radio card has difficulties recovering the data, which triggers retransmissions, excessive overhead, and low throughput. For example, a user will likely experience a 75 percent drop in throughput when operating from an area that has low signal levels.

To counter coverage holes, you need to improve the signal strength in the affected areas. Try increasing the transmit power, replacing the antennas with ones that have higher gain, or moving access points around to better cover the area. Keep coverage holes from popping up unexpectedly in the future by performing a periodic RF site survey, possibly every few months.

In a wireless LAN, a wireless client connects to an access point, and communication takes place between the client and the access point as the user browses the Internet, sends and receives e-mail, and/or talks to someone on a wireless IP phone. This communication includes an uplink path from the client to the access point and a downlink path from the access point to the client. For example, when a user opens their browser, the client device sends a URL page request through the uplink path to the access point. Then, the web pages are sent through the access point to the client over the downlink path. Another example is when a person—let's call her Mary—is talking on a wireless IP phone to another person—we'll call him Bob. In this case, Mary's voice packets flow over the uplink path from her IP phone to the access point, and Bob's voice packets

eventually travel over the downlink path from the access point to Mary's phone. At least this is what we hope occurs—otherwise, none of these applications will work properly.

Access points periodically send beacons, which travel over the downlink path from the access to the client devices. Most Wi-Fi site survey tools receive these beacons and display the signal strength and signal-to-noise ratio (SNR) associated with the beacon signals. A person performing an RF site survey determines the optimum installation location for access points by using the beacon signal strength to determine where adequate coverage is provided. For example, you might define the target range boundary to be 20dB SNR. You then use a test tool to ensure that there is at least 20dB or better SNR throughout the covered area. Keep in mind, however, that this is only in relation to the downlink path. It doesn't take into consideration the uplink signals from the client devices.

Something to consider is that most access points have a significantly higher transmit power compared to wireless clients. Access points, for example, are typically set to their highest transmit power, which may be 100mW. This is done to seemingly maximize the signal propagation and coverage from each access point in order to minimize the number and costs of access points. Wireless clients, though, tend to have a much lower effective transmit power to conserve battery power. In this situation, the downlink signal strength will be relatively high, and the uplink signals will be much weaker. This means that the effective range between the access point and the client device is governed by the uplink signal strength. As a result, the use of only access point beacons (downlink signals) for determining coverage will indicate much better coverage than what will actually be available when the clients interact with the access point. The downlink communications will be fine, but weaker uplink signals will limit the effective range and likely disrupt communications when client devices move into areas where uplink signal strength is not good enough to support communications. Figure 1-15 illustrates the uplink and downlink paths. If the client device had a higher transmit power, then the client device could go farther away from the access point and still maintain communications with the access point.

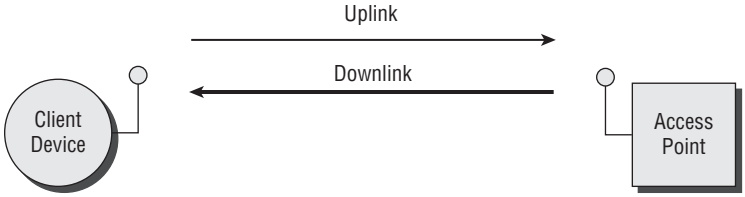

Figure 1-15: Downlink transmit power is generally greater than the uplink transmit power of client devices.

You can avoid falling into this trap with the following safeguard measures:

- **Base range measurements on the weaker uplink signal.** If the effective transmit power of the client device is lower than the access point that you plan to use in the deployment, or the signal strength of the client device at the access point (viewed by logging into the access point) is lower than the signal strength of the access point beacons at the client device (measured by the client device or signal measurement tool located next to the client device), you'll need to take the uplink signal into consideration when performing signal coverage testing. Be sure to use the weaker wireless client if the network will support multiple client devices.
- **Perform signal testing by measuring the uplink signals.** To do this, you measure a client's uplink signal strength as you walk with the device throughout the covered area. Of course you'll need to login to the access point, probably with a wireless laptop (through a web browser), to view the signal values. In addition, you'll need to periodically have the client device send something to the access point and refresh the access point display to see signal updates as you move the client device about the facility. In some cases, the weaker client device and the laptop will be the same device.
- **Consider turning down the transmit power of the access points.** If you want to maximize performance, then consider adjusting the transmit power of the access point to balance the uplink and downlink signal strengths. You can do this by turning the transmit power down to a value that makes the downlink equal to the uplink signal strength. This will increase the density of access points and increase costs due to a larger number of access points required to cover an area—but it may improve performance because of higher capacity. Fewer users will associate with each access point.

RF Interference

RF interference occupies the air medium, which delays the transmission and receipt of data and causes collisions and resulting retransmissions. These network delays may impact the operation of an 802.1X port-based authentication system. In some cases, you might need to adjust 802.1X system configuration parameters, such as timeouts, in order to compensate for impacts of RF interference (as well as other impairments).

The combination of high noise levels and high retry rates generally indicate that RF interference is impacting your wireless LAN. You can use tools such as an AirMagnet Laptop Analyzer or NetStumbler to measure noise. Also, AirMagnet

has tools for testing retry rates, and most access points store retry statistics that you can view through the Admininstration Console.

If the noise level is above -85 dBm in the band where your users are operating, then RF interference has the potential for causing poor performance. In this case, the retry rates of users will be above 10 percent, which is when users start feeling the effects. This can occur, for example, when wireless users are in the same room as a microwave oven that is operating.

If you find significant RF interference is present, then find out where it's coming from and alleviate the problem. If the symptoms occur only when the microwave oven or cordless phone is operating, then try setting the access point to a different channel. That sometimes eliminates the interference.

Also, use a tool such as NetStumbler to take a quick scan of other wireless LANs that are operating in your area. If you see that others are set to the same channel as yours, then change your network to non-conflicting channels. Keep in mind that there are only three channels (1, 6, and 11) in the 2.4 GHz band which don't conflict with each other. Most homes and small offices will have their access point set to channel 6, because that's the most common factory default channel—so you may need to avoid using channel 6 with the access points near the perimeter of your enterprise.

NOTE **In some cases, the root cause of poor performance may be an access point that has failed. Check applicable access points for broken antennas, status lights indicating fault conditions, and insufficient electrical power. Try rebooting the access points, which often resolves firmware lockups. Make sure that the firmware is up-to-date to minimize lockups in the future.**

Addressing

There are two primary levels of addressing on a computer network: IP Addressing and MAC addressing. IP addresses enable routers and the Internet to forward data packets from one point to another on the network. The use of TCP/IP-based applications requires all network devices to have an IP address. In order to go through the Internet, official and unique IP addresses must be used. For flexibility, an enterprise can use just about any IP addresses that they want to within the enterprise's network, assuming connections to the Internet are made with official IP addresses assigned to the enterprise (generally through an ISP). In order to support the assignment of unique IP addresses to all client devices, most enterprises implement DHCP, which automatically assigns and manages the IP addresses.

IP addresses correspond at Layer 3, the Network Layer, of the OSI reference model. At Layer 2, the Data Link Layer, MAC addresses must be used. The MAC address is the physical address of the NIC and is what the NIC will respond to. Every NIC has a unique MAC address assigned during manufacturing. The manufacturer's name is coded in the MAC address. In fact, most packet sniffers will display the manufacturer's name found in the MAC address of packets.

When communicating with network devices, the various IEEE 802 standards, such as IEEE 802.3 and 802.11, have protocols that send frames based on individual MAC addresses, referred to as *unicast*, or multiple MAC addresses, referred to as *multicast*. A special case of multicast is *broadcast*, where the frame is sent to all network devices.

IEEE 802.11 Multicasting

IEEE 802.1X makes use of multicast addressing, and IEEE 802.11 wireless LANs handle multicast frames in sort of an odd way that may impact performance of an 802.1X system. As a result, it's important to understand how 802.11 multicasting works. When any single wireless client associated with an access point has the 802.11 power-save mode enabled, the access point buffers all multicast frames and sends them only after the next Delivery Traffic Indication Message (DTIM) beacon, which may be every one, two, or three beacons (referred to as the *DTIM interval*). The DTIM interval can be set in the access point configuration. The 802.11 power-save mode is optionally set by the user on the wireless client device, generally via the wireless client card's configuration utility.

The use of DTIM intervals was put in the IEEE 802.11 standard to enable sleeping stations that implement the power-save function to know when to wake up and receive multicast traffic (after every DTIM beacon), and to allow flexible configuration of the network (DTIM interval) that offers a trade-off between battery life and performance. If none of the 802.11 radios associated with the access point has the power-save mode enabled, then the access point will send multicast traffic immediately and will not wait until after the next DTIM beacon.

Setting the DTIM Interval

If you haven't considered multicasting in the past, then the DTIM interval on your access points is probably set to the default value, which is likely one, two, or three. If set to one, the access point will deliver multicast frames after every beacon. Don't forget, the significance of the DTIM interval only applies if at least one wireless client associated with the access point has the 802.11

power-save mode enabled. A DTIM interval of one, two, or three is good for performance, but it will likely hurt battery life.

I've found through testing that some wireless clients with power-save enabled will receive the last multicast frame sent by the access point and then continue staying awake until after the next beacon. In other words, the client radio stays awake for the entire beacon interval. If the beacon interval is set to 100 milliseconds (the common default value) and multicast traffic is occurring within each 100 millisecond interval (typical for many multicast applications), then the flow of multicast frames will probably cause the client radios to stay awake indefinitely. This draws battery power as if the power-save mode was not enabled. Even with DTIM intervals of two or three, battery life may still suffer. For example, with DTIMs occurring every other frame (a DTIM interval of two), the radio may be awake 50 percent of the time.

Even if you have changed the DTIM interval in the past, it may still be too high. For example, I've seen some DTIM intervals set to 30, which means that wireless clients with power-save mode enabled will not receive multicast traffic until after 30 beacon intervals. This might be great for battery life, but it only allows the delivery of multicast traffic every three seconds (assuming the default 100-millisecond beacon interval). The impact of this is highly application-specific, though. For example, the transmission of voice would likely have long and annoying skips or quiet intervals, but the delivery of text messages may not experience any impairment noticeable to humans.

What DTIM interval is best? Unfortunately, there's no straightforward answer to this. In order to optimize the DTIM interval, you should gain a good understanding of how your multicast applications work. You'll probably not find this in user manuals, and even the vendor may not want to tell you much about it, to avoid disclosing company secrets. In these cases, you might need to do some testing.

Most likely, the best DTIM setting will be the one that provides maximum battery life while allowing delivery of multicast often enough to provide good performance. You should determine the highest DTIM interval that can be set and still allow the multicast applications to work properly. First, try setting the DTIM interval to one, and then use the application while monitoring the quality of the transmission (such as voice quality). In terms of DTIM settings, this should offer the best quality. You can use this as a baseline for comparison purposes when setting the DTIM interval to higher values. Now, try increasing the DTIM interval to higher values while monitoring the transmission quality. Keep increasing the DTIM interval until the quality is just above the minimum tolerable level. At this point, you'll have the highest DTIM interval that you can use to maintain required performance while also offering the best power savings.