Chapter 2

Wireless Basics

Now that I have presented a basic structure of general network concepts, I'll focus on wireless networks. In most respects, wireless networks function exactly like wired networks. They obviously use a different media—the airwaves—and follow slightly different technical protocols behind the scenes.

Although in some ways it seems mysterious that wireless networks could work at all, they rely on technologies that have been common for almost a century, going back to the invention of the wireless telegraph invented by Guglielmo Marconi in 1890. We're not surprised or concerned that an FM radio can pick up dozens of music channels across the airwaves from a distant transmission tower. We take cordless telephones for granted. Wireless networks take advantage of these established technologies to transmit computer data.

These days, I don't categorize wireless networking as a cutting-edge technology. It's been around long enough to become very stable and remarkably inexpensive, and it's gone through many generations of improvement since its debut. Just browse the shelves of your local computer store, and you'll see plenty of wireless products marketed to the non-technical home computer user.

This report favors the term "wireless network" to describe wireless local area networks based on the 802.11 family of protocols. Marketing, technical, and industry publications and Web sites often use Wi-Fi or Wireless Fidelity or WLAN (wireless local area network).

Wireless networking has a place almost anywhere computers are used—the home, small office, as well as large businesses and academic campuses. But wireless networking has really taken off in the home-computing arena. Setting up a wired network is a bit out of reach for most households; not many are willing to crawl under the house and install a wiring system.

As more households have multiple computers, as more of us do some or all of our work from home, the demand has increased, and the market has responded with many reliably functioning products remarkably easy to set up and use. Affordable and easy-to-use wireless equipment makes having a full-fledged network in the home a viable possibility.

Wireless Evolution

Wireless has evolved to the point of being a commodity technology, which means you don't have to worry much about whether or not the equipment will work as advertised. With millions of units sold, most of the kinks have been worked out by now.

This wasn't the case in the early days of wireless networks, when the equipment was very finicky and you had to do a lot of testing and adjusting to end up with a functional wireless network. We're past that now; it's pretty much a plug-and-play world—thank goodness! The equipment is developed to the point that just about anyone can set up a wireless network at home.

One of the most popular consumer devices is the wireless router. Designed to provide a broadband Internet connection for multiple computers in the house (or neighborhood!), a wireless router is hooked up to a DSL or cable modem connection. A wireless router bundles a number of functions into a single device. On one side it must have a cable modem or DSL interface. It functions as a wireless access point, a DHCP server, and it creates a private IP network through a built-in IP router with network address translation (NAT). These devices are ideal for the home or small office but wouldn't be recommended for larger business networks.

802.11 Media Access Rules

Wireless networks follow most of the same media access rules as traditional Ethernet–with one important difference. They share the CSMA rules but operate on a principle of Collision Avoidance instead of Collision Detection. In a wireless environment, dealing with collisions is more complex than with optical or copper media. On a wireless network, there may be some devices out of reach of others, making it impossible to detect collisions. Collisions, therefore, must be avoided.

The media access rules of 802.11 include the use of Request to Send (RTS) and Clear to Send (CTS) signals, which ensure each transmission happens when the media is idle, avoiding most collisions. Before a station transmits, it broadcasts an RTS signal and issues a CTS signal when finished. The more orderly network conduct of collision avoidance, although seemingly more polite, results in more overhead and less efficiency than the collision avoidance strategy of the Ethernet world.

Wireless Architectures

Let's take a deeper look into the structure of wireless networks. They can be organized in two basic ways: ad hoc mode and infrastructure mode.

- **Ad hoc mode** involves individual computers equipped with wireless cards communicating directly with each other. This method isn't widely used but is available for sharing files from one computer to another or using one computer to print to another's printer.
- **Infrastructure mode** stands as the more common wireless architecture. This approach involves wireless-enabled computers connecting to a wireless access point, which, in turn, usually connects to a wired network. In this arrangement, the computers inherit access to the resources of the wired network, including access to the Internet and to any other servers and services available. *Infrastructure mode is the basis of all business and academic wireless networks*.

Most users will rarely, if ever, need to use ad hoc mode. Many wireless drivers can be set to ignore ad hoc wireless networks, because these networks usually consist of individual computers that don't accept connections and have no services to offer.

With ad hoc mode, it's possible to allow multiple computers to connect and share files and printers in the absence of additional infrastructure. Also called "peer-to-peer" mode, this approach requires the Wireless NIC (network interface card) of two or more computers to connect directly to each other. Often, when viewing the list of available wireless networks, you can see other computers configured to display their SSID (session set identifier or identification) in ad hoc mode. Rarely, however, are these connections useful for connecting to the Internet. Applications of ad hoc mode are usually limited to sharing files and printers, hopefully with adequate security enabled. One wouldn't typically use this approach for Internet access–that's much more easily accomplished using infrastructure mode, making use of access points or other devices.

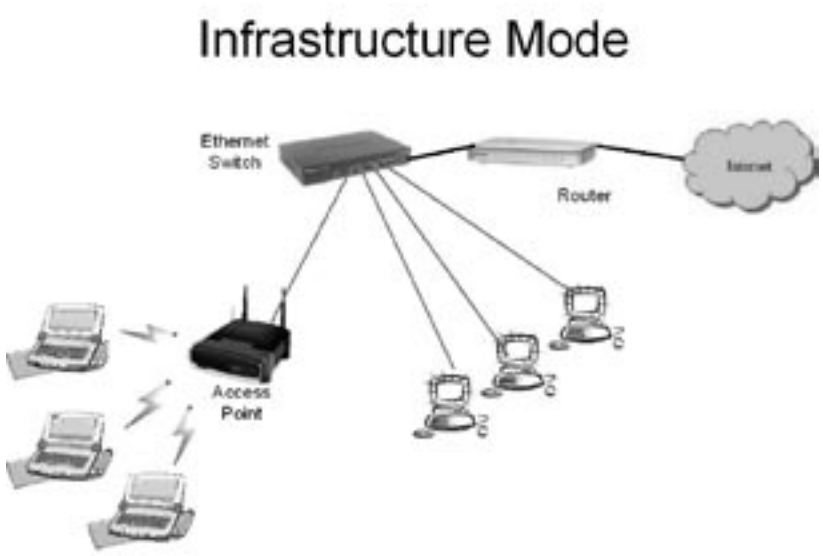

Figure 3

By far, the most common approach for setting up a wireless network is through Infrastructure Mode. This is the approach that I'll be talking about for the remainder of this report.

As Figure 3 illustrates, a wireless network in its simplest form consists of a wireless access point connected to a wired network, which offers connections to individual computers equipped with wireless network interface cards. Only rarely does a wireless network exist in isolation. They almost always connect to a wired network.

Wireless networks should be compatible with all the types of computers involved. It's acceptable to mix laptops and

desktops, Macs and PCs, all regardless of brand. It's usually fine to mix brands and models of wireless cards and access points, too. As long as all the equipment conforms to the standards, which I'll discuss in more detail later, all the various brands and device types will work together.

Wireless Hardware

Let's take a close look at some of the devices involved.

A **wireless access point** serves as the primary piece of communication equipment. It functions much like the Ethernet hub described previously. Instead of having RJ-45 jacks to connect physically to computers, it has a set of radio frequency (RF) transmitters and receivers. (I'll discuss the details of RF communications in upcoming sections "RF Interference" and "RF Modulation Schemes.") The access point will have a single Ethernet jack, which is used to connect it to the wired network. While some access points can obtain the power they need from the Ethernet connection (power on Ethernet) most will require an AC outlet.

In that it is a shared media, an access point shares another important characteristic of the Ethernet hub. In the same way that all the devices plugged into an Ethernet hub have technical access to the packets transmitted to and from the other devices that connect to the unit, the same is true for the access point. This isn't an ideal situation from a privacy and security vantage point. I'll talk about these issues more later on, but keep this important point in mind.

End-user devices, such as notebook or desktop computers and PDAs, must be equipped with a wireless network interface card (NIC) in order to participate in a wireless network. While most notebook computers sold today have wireless capabilities built in, older units will need a wireless NIC added *a la carte*. Desktop computers typically don't have built-in wireless because most are stationary and would more likely be connected to a wired network. There are wireless cards available for Macs, PCs, and most PDAs. In addition to PCI and PC Card (formerly PCMCIA) interfaces, USB (universal serial bus) is used for many wireless NICs.

Transmission Details

One essential component of a wireless access point involves its radio frequency transmitter and receiver. Both access points and wireless NICs must have the capability to transmit and receive on the frequencies specified by 802.11. While 802.11b and 802.11g operate on the 2.4 GHz spectrum, the transmissions actually take place on several discrete channels within that broader range. (For descriptions of the differences among the various 802.11 designations, refer to the upcoming section "Wireless Flavors.") The standard specifies fourteen channels, but the use of the channels varies according to country-specific regulations. In the United States, Federal Communications Commission (FCC) rules allow the use of eleven channels. As shown in Figure 4, each of these channels consumes 5 MHz of the allocated spectrum and overlap.

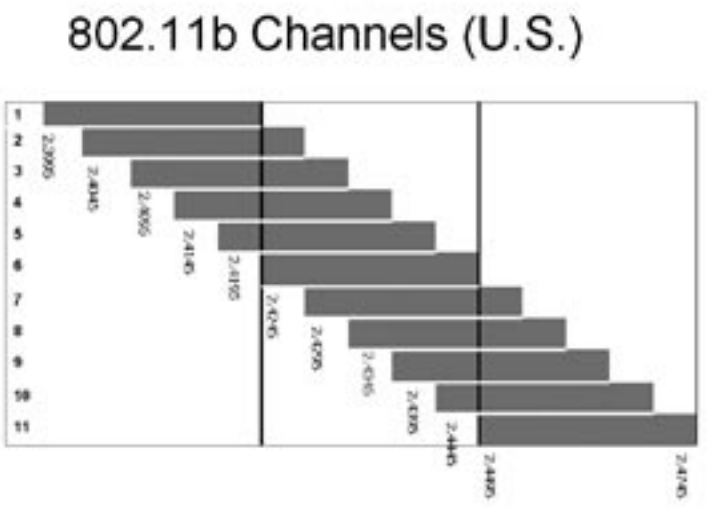

Figure 4
Transmission Channels defined in 802.11b

One of the important considerations in a wireless network design involves channel selection. In a simple WLAN (wireless local area network) with only one access point, any of the available channels can be selected. In an environment with multiple access points, it's necessary that adjacent access points be set to different channels. The channels selected should be separated enough to ensure that they don't overlap. This requirement effectively limits the number of channels that can be used in practice to about three. For wireless networks in the United States, channels one, six, and eleven might be selected to ensure non-overlapping transmission. (Figure 4 illustrates the eleven channels allocated for use in the United States and the fact that only three of these do not overlap.)

Network access cards do not have to be configured to select a channel—they are able to detect and negotiate the channel associated with the nearest access point.

RF Interference

One of the problems with 802.11b and 802.11g and use of the 2.4 GHz band stems from the possibility of interference from other devices. Several other types of equipment also operate on this frequency, including microwave ovens and cordless telephones. (Cellular telephones operate on a different part of the spectrum and do not interfere.) Given the strength of most microwave ovens, one should be sure to locate any access point at least ten feet away from them. Industrial microwave ovens may operate at even higher power levels, and access points should be positioned accordingly. If you see a pattern of poor performance and dropped connections around lunchtime every day, you might find out that placing the access point on top of the microwave oven in the staff break room wasn't such a good idea.

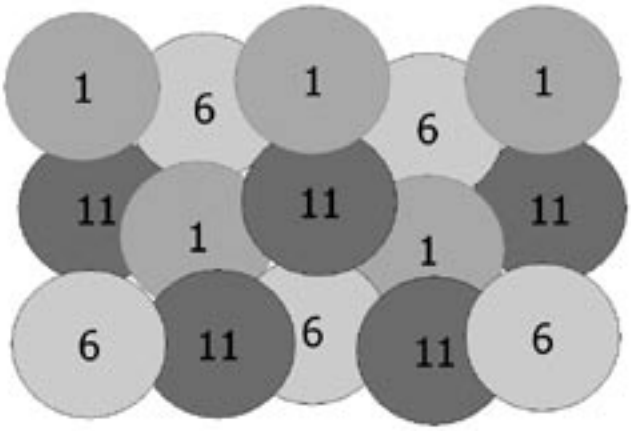

Figure 5
Non-Overlapping Channel Map

Bluetooth also operates in the 2.4 GHz space but only rarely interferes with WLAN. Bluetooth relies on a high-speed frequency hopping technology that, even though it operates in the 2.4 GHz range, only under the most unusual of conditions poses a problem with interference to 802.11b or 802.11g networks.

In most cases interference from other consumer devices will result in only a minor decrease in performance, but the placement of access points should be made in consideration of these devices. When performing a site survey, it's important to be sure that all known radio frequency devices are operating so the best position of the access points can be determined.

RF Modulation Schemes

802.11b relies on an RF transmission modulation scheme called "direct sequence spread spectrum," or DSSS. Rather than transmitting on a single frequency, 802.11b disperses the signal over a 30MHz range within the 2.4 GHz bandwidth. This technique of spread spectrum makes transmissions less susceptible to interference. The width of the spread means there are fewer overlapping channels available for practical use.

DSSS contrasts with "frequency hopping spread spectrum," or FHSS, a scheme in which the carrier signal frequency changes several times per second. The transmitter and receiver are programmed with an algorithm that allows them to stay synchronized as they rapidly hop through a pseudorandom sequence of channels. Frequency hopping makes it difficult to jam or eavesdrop on transmissions and results in minimal interference from devices that might share its part of the spectrum. Bluetooth is based on FHSS.

Positioning Wireless Access Points

If you intend to cover a large area with your WLAN, you'll need multiple access points. These access points will need to be positioned so they provide overlapping areas of coverage, allowing users to roam throughout the hotspot without encountering dead spots. In a completely open environment, the overall WLAN could be pictured as overlapping circles of coverage with an access point at the center of each circle (see Figure 5).

The real world is more complicated. Dense objects can block the signal, producing a dead spot in the area behind it. Multi-floor buildings can be challenging as well. An access point positioned in a ceiling can often provide coverage for two levels of the building. An access point near an exterior wall of the building can provide some outdoor coverage. Corporate organizations or academic campuses may want to have wireless coverage between buildings. In come cases, it's important the wireless network not extend beyond the controlled space of its buildings (see "Eliminating Rogue Access Points, Chapter III, p. 22).

One of the major factors in positioning access points involves the user density in each area. In areas where a high number of users may need to access the wireless network, it may be necessary to deploy more access points than would be dictated by geographical factors alone. Keep in mind that access points are shared media devices, and a large number of simultaneous users with continuous bandwidth demands may exceed a single access point's capacity.

Remember, the RF transmission that underlies wireless technologies is not a "line-of-sight" technology.

In the same way that you don't expect to see the transmitter tower that broadcasts the FM radio station you receive on your car radio, an access point generally isn't visible to the clients that use it. The radio waves easily penetrate walls and other solid objects. It's possible to use an access point to provide wireless coverage throughout several rooms in a building.

Also keep in mind that there are some materials that absorb or reflect the signal. Substances that contain a high water content level tend to absorb some of a 2.4 MHz signal. Metal objects cause signal reflections that may create dead spots.

For a simple wireless network, it's possible to use an informal approach in setting it up. Particularly in the home or small office—which might be well covered by a single access point and may have relatively few interfering devices or obstructing materials—one can pick a central location, hope for the best, and probably have a functional hotspot. If the space to be covered is large and complex, then a professional RF site survey should be performed. Most computing consulting firms will have an engineer with the equipment and training necessary to perform a site survey for a wireless network. Organizations with a large information technology department may have this capability in-house.

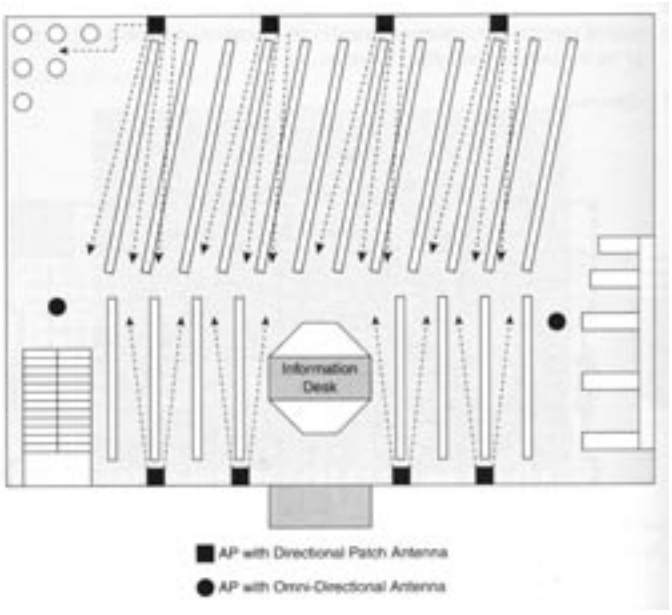

Figure 6
Sample Library Site Map, from *802.11 Wireless Network Site Surveying and Installation* by Bruce Alexander, reproduced by permission of Cisco Press

A professionally conducted site survey will lead to a WLAN design that covers the geographic area desired, has no dead spots, supports the number of users required within each cell, and uses the fewest number of access points possible. Some of these goals can be achieved by over-saturating the area with access points, but this approach significantly increases the cost of the WLAN.

In some circumstances, commissioning a professional site survey is a prerequisite to obtaining funding for a WLAN project. Having the site survey accomplished beforehand, too, is often necessary to determine the costs of the equipment required. The network documentation (produced via such a site survey) also will give the funding source confidence the project will be successful.

Library buildings can be especially complex spaces for installation of wireless LANs (see Figure 6). While reading rooms tend to be open and Wi-Fi friendly, book stack areas are notoriously difficult areas in which to provide uniform wireless access. Books have a high density and high moisture content and tend to absorb RF signals. Metal shelves cause reflections. Given that it may take a large number of access points to cover the library's stack areas and the lower likelihood that library patrons will take advantage of wireless access in these areas, many libraries choose not to include them or live with spottier coverage.

Once you've planned the location for your access points, you can be creative about how to physically mount them. It's often a good idea to position points out of plain view. Not only might they detract from aesthetics, especially in public areas of the library, but having them out of sight also will reduce chances of theft and tampering.

It's common to mount access points in the space above ceiling tiles. If the area above the ceiling is part of the building's air circulation system, however, it may require plenum-rated equipment and cabling. Your building may have cabinets, panel cut-outs, or other spots well suited for access points. Wall mounting is also a popular option (see Figure 7).

Remember, too, you'll need an Ethernet drop and electrical outlets installed. The placement of the access points is often influenced by locations where it's practical to install network and power cables.

Range per Access Point

The amount of area that can be covered by a single access point depends on the type of equipment (802.11b, 802.11a, or 802.11g) and the physical characteristics of the space and the density of materials located throughout the space. That said, there are some general expectations. Indoors, a single access point usually can cover a radius of 75 to 100 feet in a typical office configuration. In large open areas, coverage may extend to as much as 500 feet. Performance will be better for those users situated nearer to the access point (compared to those at the outer edge of coverage). The range for ideal connectivity will be much smaller than the range that achieves a partial level

of performance. 802.11a networks have a shorter range than those based on 802.11b or 802.11g.

For outdoor networks, greater range can be accomplished using access points designed for outdoor deployment with more powerful antennas as well as with housings designed to endure exposure to weather conditions.

Antenna Options

Wireless access points and network interface cards come with built-in antennas designed to work well for most conditions. Some wireless devices have antennas that cannot be replaced. Many laptop models have the wireless circuitry built into their motherboards with antennae integrated into the cases. Most access points come with built-in antennas, but these can be replaced to accommodate special circumstances.

The most popular antennas for wireless networks are omni-directional, meaning they transmit and receive equally well in all directions. You can visualize the range of this antenna type as a sphere, though some have less vertical coverage than horizontal.

Directional antennas can be used when more specific targeting of the signal is required. These antennas can be used to support point-to-point wireless links or when the coverage area is more elongated than circular. Directional antennas often transmit at higher power levels. Some of the types of directional antennas include the conventional dish style as well as Yagi, an antenna with a long, rod-shaped design.

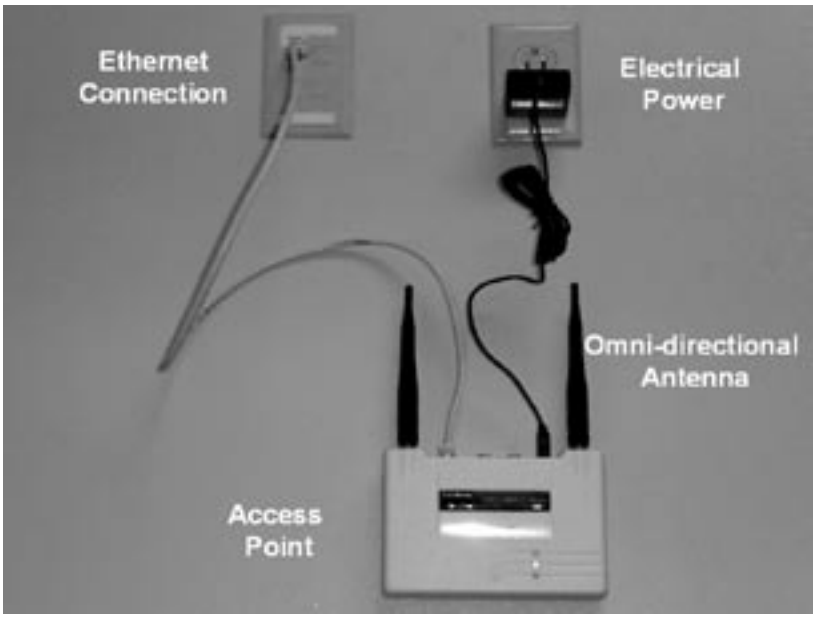

Figure 7
Wall-Mounted Access Point

Wireless Equipment Types

Although it's possible to use ad hoc mode to connect end-user devices with wireless cards directly to each other, by far the most common approach involves using a central communication device.

The most common device, called a "**wireless access point**," establishes a hotspot of wireless connectivity. In most cases the wireless access point will connect to a wired network, making all the services on that wired network available to wireless users. An access point includes a number of internal hardware and software components.

In one respect, the access point functions as a bridge, connecting two networks with different media access protocols. The wireless access point bridges an 802.3-based network (governed by the CSMA/CD media access rules) with the 802.11 rules, which are governed by CSMD/CA (see page 14, "802.11 Media Access Rules").

The access point also provides the same function to wireless devices as the shared media hub does on an Ethernet LAN. Multiple end-user wireless devices connect through a single access point. They share the media provided by the access point, competing for the overall bandwidth available and contending with each other for collisions.

Physically, the access point includes a radio frequency transmitter and receiver tuned to operate on the 2.4 GHz band. You can usually see the radio antennas associated with this part of the device. The antennas on most access points are non-directional, producing a circular hotspot with the AP (access point) in the middle. Directional antennas can be used to cover irregular spaces when needed.

The access point will have an Ethernet port in the form of an RJ-45 jack. This port allows the AP to connect to an Ethernet network using a standard cable. The Ethernet port of the access point can connect to an Ethernet hub, a switch, or a router.

Most access points have a power adapter that connects to a standard AC outlet. While some models are able to tap the power they need from the Ethernet connection using a method called "Power-over-Ethernet" (PoE), most inexpensive models require a separate power supply. PoE requires additional equipment on the head end that infuses power onto an unused pair of wires in the Ethernet cable.

Another type of base station, the **wireless router**, also called a "wireless gateway," includes all the functions of the access point, plus the ability to share an Internet connection on a

small network. Most business networks, including those in libraries, have the infrastructure already in place to provide Internet access to all the devices on the network. In most cases, a router has been installed and configured to deliver packets to and from the Internet. Most homes and many small offices, however, lack this infrastructure. A wireless gateway makes it possible to share a high-speed DSL or cable modem Internet connection among a number of computing devices equipped with standard wireless cards.

The wireless router combines a number of components into a single unit that, from its external case, often looks no different than an access point. The physical layout of the wireless router mirrors the access point: antennae, RF transmitters/receivers, Ethernet jack, and power plug. Internally, these units have additional components that produce a complete self-contained network.

As implied by the name, this equipment contains a full-fledged router, capable of managing the transfer of data between end-user devices and the Internet. The device establishes a LAN from the devices under its control, taking care of such details as assigning IP addresses. Most include the ability to assign addresses according to a private IP network using a built-in DHCP service. Wireless routers perform Network Address Translation (NAT) so these private addresses can communicate effectively on the Internet.

Optionally, a wireless router may have built-in connections for DSL or a cable modem, eliminating the need for external cable or DSL modems. Some of these devices also include a set of Ethernet ports allowing wired devices to share the Internet connection.

Access point or router? Keep in mind that wireless routers mainly are designed for the home and small-office environment so multiple users can share a high-speed Internet connection. Access points are used to extend an existing LAN. If you already have a LAN with Internet access, you need an access point, not a wireless router.

Wireless Flavors

Wireless networking specifications have gone through a number of iterations. The IEEE 802.11 committee worked through many issues having to do with the underlying transmission technologies, eventually developing a specification (**802.11**) that operated at 1 to 2 mb/sec. Although this groundwork was important, wireless networks based on the original specification did not have widespread deployment.

802.11b–The first version wireless networking that caught on in the wireless networking arena was 802.11b. This version offers data transmission speeds of up to 11 mb/sec, operating on the 2.4 GHz part of the spectrum. Products based on 802.11b began shipping in 1999, with sales growing through about 2003. Although 11mb/sec is the theoretical maximum, actual performance tends to range in the 2-4 mb/sec range.

This was a very popular technology that gained wide acceptance in the business and home consumer arenas. Many millions of units of 802.11b equipment are still in service, but sales have diminished. Though one would likely purchase Wi-Fi equipment based on newer protocols, the huge installed base of 802.11b equipment makes compatibility a vital consideration.

802.11a–By today's standards, the 11mb/sec possible with 802.11b is rather slow, especially given that the actual throughput is substantially less. Fortunately, wireless technologies have evolved to offer faster performance. 802.11a delivers faster performance with a maximum throughput of 54 mb/sec. This flavor of Wi-Fi operates on the 5 GHz band. While this part of the spectrum suffers less from interference from other devices, it's not compatible with the large installed base of 802.11b equipment. It also has shorter range, with the maximum distance between the client and the access point limited to about 180 feet.

The advantages of 802.11a include higher performance and less interference. Disadvantages include incompatibility with the large installed base of 802.11b equipment, shorter range, and higher cost. This version of wireless networking tends to be used in closed business environments, not in public hotspots.

802.11g builds on the technologies present in 802.11b. Operating on the same 2.4GHz band, 802.11g delivers a maximum of 54 mb/sec. Since it's based on the same transmission technology, it's easy to build equipment that's backwardly compatible with the existing 802.11b equipment. Most access points designed for 802.11g can support both 802.11b and g clients. The b clients will operate at the maximum 11 mb/sec throughput, however. A computer equipped with a g network interface card can connect to a b access point but will step down to 11mb/sec.

Just as with the other wireless flavors, the actual throughput achieved with 802.11g will be less than the theoretical maximum; the fastest data transfer rates will likely be no faster than 20-25 mb/sec. Users at the outer extent of the range will see performance diminished even further.

Because it operates in the 2.4GHz range, 802.11g is subject to the same RF interference as the earlier 802.11b version. The range in which one can achieve anything close to maximum throughput is shorter than that needed to maintain 11mb/sec in an 802.11b network. To uniformly achieve high performance in an 802.11g network, a network will need a larger number of access points positioned closer to each other.

802.11g is able to achieve higher data rates through Orthogonal Frequency Division Multiplexing (OFDM), a transmission scheme that divides the signal into multiple chunks that are then simultaneously transmitted at different frequencies. 802.11a also takes advantage of OFDM to deliver fast performance.

At the time of writing, most new equipment sold for home and small office use is based on 802.11g. Unless you have very specific needs for which 802.11a might be preferred, most new library wireless networks should be based on 802.11g. Those based on 802.11b are not necessarily considered obsolete, but components acquired to expand the network or replace failed units should be based on the newer standard. In some cases, a more wholesale replacement of 802.11b may be needed if there is great demand for higher performance.

802.11n (future)–Today, the fastest wireless LANs operate at 54 mb/sec. Although that performance level seems much better than the previous generation of 11mb/sec, it's still slow relative to what's possible with wired networks. Work is underway to develop a wireless networking technology that will deliver data throughput of 100 mb/sec. Multiple proposals on how to best accomplish this performance boost are still under consideration by the IEEE 802.11 Task Group N, which is charged with defining the standard. It's expected the standard may be completed as soon as the end of 2005. Products supporting this standard may be about two years away.

Wired vs. Wireless Bandwidth

Wireless technologies continue to increase in performance, but wired networks will always be able to deliver higher performance. Still, wireless networks offer mobility and convenience.

Consider the performance available via wired networks. Today, 100 mb/sec Ethernet to the desktop is quite common. Very few Ethernets of 10 mb/sec remain in use. A growing number of the 100 mb/sec Ethernets are based on switches, fewer rely on shared media hubs, and the performance capacities for wired networks are continuing to increase at a rapid pace. Gigabit Ethernet has been available for several years and is implemented routinely on servers and backbone network linkages. The cost for gigabit Ethernet equipment has dropped considerably, and many organizations use it for standard desktop connections. Gigabit networks cards for a PC sell for as little as $30 and low-end gigabit switches for less than $100. At least for the home and small-office environment, gigabit Ethernet is the current commodity-level technology.

On the higher end of the wired networks, 10 gigabit equipment is currently available and is deployed for high-performance servers. Currently, 10Gb Ethernet is primarily implemented with fiber optic cabling. A version compatible with high-grade unshielded twisted pair cabling is expected in 2006.

On a gigabit Ethernet, one can see file transfer speeds around 750 mb/sec. Compare this figure to the 25 mb/sec throughput available on the fastest 802.11a or g wireless network, and it's clear the typical wired network is about thirty times faster. The performance increases for wired networks have evolved faster than increases for wireless. Each new generation of wired Ethernet tends to offer a ten-fold improvement. The evolution from 802.11b to 802.11g yielded only a five-fold advance; 802.11n also promises a five-fold boost.

One of the key trends in computing deals with the increased interest in applications that consume huge amounts of bandwidth. Sound and video regularly complement text, and scientific and educational computing involves ever-larger data sets. Even gaming applications may have significant network bandwidth requirements. Given this increased need for high-performance networks, most organizations will continue to need wired networks for servers and stationary desktop computers.

Because of its performance limitations, wireless will not displace most wired networks; yet it fills an important niche in supporting mobile computing. Those with the opportunity to build a new library, for example, should not be tempted to avoid installing a high-performance wired network. It takes *both* wired and wireless networks to meet all the needs of a modern organization.

New building construction should always include the installation of the highest quality network cabling possible, with connections provided to all possible desktop computing locations. One also needs to be sure that wired connections are provided to all possible locations for wireless access points.

Still, the idea of the "all wireless library" is often mentioned as a technology trend. One must be careful to note this has more to do with providing wireless coverage throughout the building rather than eliminating the use of wired networks.