CHAPTER 9

Wireless LANs in the Enterprise

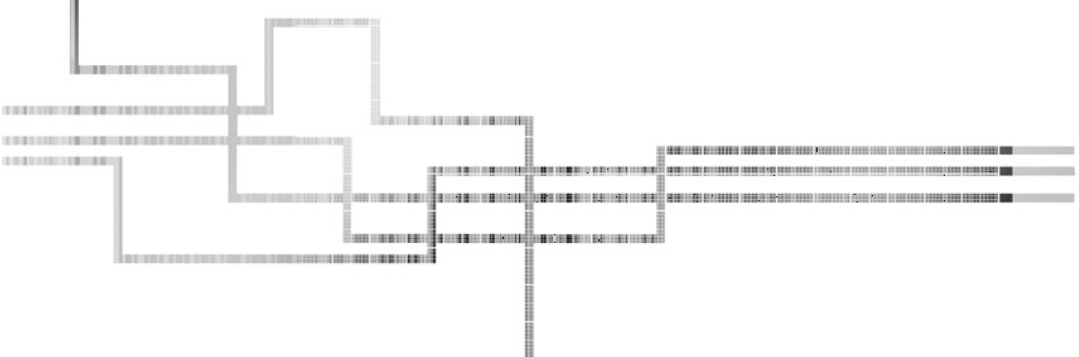

As discussed in Chapter 1, early forms of wireless LANs have been available since the mid-1980s. Standardized and interoperable WLANs have been shipping since 1997, and Wi-Fi products have been available since 1999. Despite this and the more than $1 billion spent worldwide annually on 802.11 products at this writing, WLAN deployments in the enterprise are still in their earliest days. Today, to the extent that WLANs are found in the enterprise, they tend to be limited deployments in places like conference rooms, cafeterias, and, naturally, the senior executive floors. Indeed, WLAN proliferation into vertical markets like retail, manufacturing, and warehousing greatly exceeds enterprise adoption. As evidenced by the sheer number of low-cost, easy-to-install Wi-Fi products available at computer retailers and catalogers, proliferation of wireless into residences and small offices is growing rapidly—in fact, far more quickly than into enterprises. Today's enterprise deployments are almost experimental in nature, as enterprise IS (information services) professionals, managers, and staff gain familiarity with WLANs and come to understand how they can best integrate Wi-Fi into an overall enterprise information infrastructure.

In this chapter, we define enterprise WLAN deployments, making, at the functional level, a distinction between enterprise deployments and small office/home office (SOHO) deployments. We discuss the approach an enterprise typically takes when deploying a WLAN. We also outline the steps enterprise IS professionals should take to maximize the likelihood of a successful initial deployment, including a physical assessment of the facilities in which wireless is to be deployed (the site survey) and the capacity planning needed to provide the enterprise-level performance demanded by users. Given that in the enterprise an existing wired LAN already exists, we discuss how IS professionals best can integrate Wi-Fi into this overall infrastructure, where wireless adds a vital mobility element to a network and where it might be a replacement or alternative to more traditional wires. We also discuss how you can best leverage existing network management tools and practices from the wired world to most expediently bring a similar level of management to the WLAN.

A theme throughout this chapter is that 802.11 equipment should be considered a highly integrated network element, rather than simply tacked onto a LAN, whether the network is in the home office or resides within a large enterprise.

WHAT IS THE ENTERPRISE?

First and foremost, the *Enterprise* is, of course, a series of starships, all captained by dashing leaders and crewed by a pan-galactic collection of Federation officers. Having said that, a definition of the enterprise as it relates to WLAN deployments is probably more germane to this book. Like the *Enterprise,* many enterprises are large, consisting of, at minimum, hundreds of individuals, all of whom are users of the organization's information infrastructure in some fashion. While many of these users may be located in a single headquarters building or campus, the enterprise is typically geographically distributed, with users scattered across a region, a continent, or even around the world.

The fact that a user may be working out of a spare bedroom thousands of miles from the enterprise headquarters makes that user no less an enterprise user—indeed, it is these sorts of users who often most challenge IS professionals.

NOTE The average 802.11 sale to enterprises consists of three to five access points, because most enterprises worldwide are small businesses rather than the more widely publicized large corporations.

Certainly, large commercial entities around the world are considered to be enterprise-level organizations. The more expansive definition is one that includes any large organization with a common purpose where individuals are engaged in specific, complementary tasks—including managing the enterprise information infrastructure. By this definition, larger governmental entities on the city, county, state, and provincial level as well as on the national level are "enterprises." Similarly, school systems, whether public or private, are enterprise organizations.

Stated another way, an enterprise is any organization that reaches the size at which it requires a dedicated staff of one or more IS professionals. The charter of this staff, no matter how small, is to make certain that the information infrastructure meets the needs of the organization and enables it to meet its goals—ideally, better than competitive organizations. And to remain *competitive,* leveraging new technologies to its advantage is a requirement for any organization, whether in the private or public sector. Wireless LANs are an excellent example—perhaps the best example available today—of an information technology that can have dramatic impact on the efficiency and effectiveness of an organization. Not surprisingly, IS professionals around the world are increasingly being charged with installing Wi-Fi, often on a trial basis with small pilot programs but with a mind toward a ubiquitous enterprise-wide deployment.

A SOHO Wi-Fi deployment presents few of the challenges associated with an enterprise deployment. As discussed in Chapters 5 and 6, the range of Wi-Fi devices, varying from a low of 60 feet to over hundreds of feet, is more than sufficient to cover even the largest of homes and small offices—even at the highest possible data rates. Indeed, if a home is so large as to require more than a single access point to achieve full physical coverage, it's likely that the owner has the wherewithal to hire an IS professional to manage the installation. The number of users in a SOHO environment tends to be fairly limited. While it's true that users of a SOHO LAN use the LAN to access other local computers, it is far more typical for users on a smaller LAN to access data from across the WAN, which can be cable, DSL, or even dial-up. This sets their performance expectations at fairly modest levels.

The enterprise is, of course, a completely different story. Typical corporate, governmental, and educational facilities, by their multistory nature alone, require more than a single access point to cover the entire building. In campus settings, the requirement can even be extended to include not just complete and reliable in-building coverage, but also WLAN availability between buildings. This opens up a whole host of challenges not found in SOHO deployments including *roaming* and *channel reuse* that will be discussed further in this chapter.

In larger enterprises, users have come to expect a level of network performance that is consistent with a wired network, one that is often switched, providing dedicated bandwidth that is typically rated at 100Mbps—and occasionally faster. Given this level of expectation, the IS professional's challenge is to provide the freedom and flexibility of wireless with performance and security that approximates that of the wired network.

In short, an enterprise is a relatively large organization with a common goal. The organization is typically in some form of competition with organizations with similar goals and, as such, employs information technology (among other tools) to gain competitive advantage. As such, the deployment of Wi-Fi in an enterprise presents challenges not found in other sorts of deployments and substantial consequences when things don't go quite as planned.

WI-FI DEPLOYMENT IN THE ENTERPRISE

Like any large project, the first step is to set goals and then formulate a plan to meet those goals. Although the specific goals of an enterprise's Wi-Fi deployment will vary, there is a constant: to deploy a Wi-Fi network in designated areas that provides reliable coverage and delivers the expected level of performance without compromising corporate security. Although this sounds simple, as the saying goes, "The devil is in the details."

Designating Areas

Rare are the cases in which a large enterprise chooses to deploy a Wi-Fi network across the whole organization in one fell swoop from initial deployment. There are a few reasons for this. Obviously, finding the budget for what can be a significant financial undertaking can be quite difficult. Responsible financial planners tend to take more of a "show me" approach, requesting first that a pilot program be run to assess the expense and resource drains of the project, the veracity of the budget estimates, and the return on investment.

Additionally, IS professionals recognize that Wi-Fi has a learning curve (as is typical with any new technology), and running a limited deployment provides valuable on-the-job training. Finally, as was discussed in Chapter 6, Wi-Fi is a technology undergoing rapid change, and organizations have concerns, unfounded or not, that the product they deploy will lock them into a soon-to-be-obsolete technology.

The great majority of enterprises instead initially opt for a limited WLAN deployment. There are different criteria by which these deployments can be limited, as described in the following sections.

Limiting Deployment to Only Where It's Needed Most

This strategy is based on the assumption that when laptop users are in their base area, such as an office, cubicle, or desk, they access the network via a wired connection, either by plugging directly into an Ethernet jack or through a docking station. Therefore, the Wi-Fi deployment is limited to places people tend to congregate *away from their*

desks, in areas like conference rooms and smaller meeting rooms, cafeterias, classrooms, auditoriums, lobbies, and other similar public areas. For many enterprise organizations, this strategy meets the "80-20 rule"—it deploys WLANs in the 20 percent of places where 80 percent of it will be demanded.

What this strategy doesn't take into account is the fact that people are unpredictable and the places where they meet to collaborate are not always where the building's architect envisioned. Information is exchanged (and required from the network) in a variety of places: leaning up against a coworker's cubicle, in the smoking "lounge" outside, in the hallway. . .wherever.

This unpredictability continues to increase as enterprises more commonly issue laptops (as opposed to desktop computers) and as more than data is being transferred. Also, the growing popularity of personal digital assistants (PDAs) and devices such as bar code scanners and 802.11 handheld phones, and the associated demand for them to be just as connected as a laptop, drives demand for a more ubiquitous wireless infrastructure, because people use PDAs and similar devices in more places than they would a full laptop.

Similarly, as organizations begin to use the Wi-Fi infrastructure to provide local voice support, the user expectation is that coverage will be as complete as for their cellular telephone—only more reliable. For other organizations, deploying Wi-Fi "only" in the classrooms and auditoriums is tantamount to a full deployment. If a limited deployment in kindergarten through twelfth-grade schools, colleges, and universities is desired, another means of limiting the deployment is necessary—leading us to the next strategy.

Limiting Deployment to One Building at a Time

In campus environments, particularly those campuses where different buildings or groups of buildings have differing charters, it's common for Wi-Fi to be rolled out on a building-by-building basis. This is a very typical model in a university where, for example, the business school deploys WLANs in its building and then supplies Wi-Fi client adapters to (or mandates their purchase by) all students who use that facility.

Often, the financial structure of a university plays a role in the choice of this strategy. Using the business school example again, the business school may have the budget autonomy to fund an initiative to deploy Wi-Fi without the involvement of the university's central organization, and may be able to rely on outside sources of funds such as alumni associations and local business partnerships.

Sometimes, a single-building deployment is accomplished even without the involvement of the central IS organization, although this is more common, not surprisingly, in an engineering school than in a business school. The central drawback to this approach is that all but a few students and even some faculty spend their academic days in more than a single building or group of buildings. This is all the more true of new matriculates—the very ones who are receiving the first client adapters.

Experience has shown that once WLANs are deployed in a single building, the expectation is set that it should be similarly deployed across campus, in classrooms,

cafeterias, and unions, and even in the dormitories. As discussed in the next section, this unmet demand can have very real ramifications for the whole of the IS infrastructure.

Limiting Deployment to Temporary Buildings and Workgroups

In this model, Wi-Fi is deployed not so much for the mobility it provides the user, but rather for the mobility it provides the *infrastructure*. In today's dynamic economic environment, it's common for organizations to rapidly increase and decrease in size. It's also common for groups of people from different groups and even locations to be brought together on a temporary basis for a specific project. This phenomenon has fostered the creation of the term *networks in motion.* Enterprise organizations sometimes deploy a Wi-Fi network to meet these challenges.

With a temporary building, there's little economic sense to installing Ethernet cable throughout a building, or the far more expensive option of trenching for either Ethernet or fiber optics, only to soon leave it behind. Often, a temporary building has a copper plant in place that supports a telephone system, but the cabling is insufficient for modern information networks. Temporary cabling solutions with cable exposed hanging from ceilings, between buildings, or duct taped to walls present an unprofessional appearance and potential safety hazard inconsistent with most enterprise organizations' standards. A Wi-Fi network can be deployed far more rapidly throughout a building than a traditional network infrastructure and with far less expense. When it's time to vacate the building, the network infrastructure can be easily packed up and redeployed at the next location.

Temporary workgroups present challenges similar to those of a temporary building and are similarly well suited to a Wi-Fi deployment. Again, Wi-Fi networks can be rapidly deployed in areas like cafeterias, gymnasiums, tents, and the like that are designated for a temporary workgroup, including emergency or disaster relief organizations, or for business continuity purposes in the event of a local disaster. WLAN deployments greatly mitigate the "spaghetti problem" of Ethernet cable being run to individual workstations. Wi-Fi equipment can easily be easily deployed—and redeployed.

It is these sorts of installations in the enterprise that drive a significant portion of the demand for client form factors, such as USB and PCI, that are designed for desktop, rather than laptop, PCs. Industry data shows that these form factors account for as much as a quarter of all client adapter unit sales, suggesting that the deployment of Wi-Fi LANs for temporary buildings and workgroups is more common than is immediately intuitive.

Limiting Deployment from the Outside In

Enterprise organizations report that, on average, around 30 percent of all branch offices and/or their personnel will relocate over the course of a single year. This presents major challenges for enterprise IS staff—handling network additions and moves is a costly and time-consuming exercise in any event, but performing them on a remote basis in a branch office presents an even greater challenge.

While the smaller size of a branch office tends to decrease the need for the mobility Wi-Fi provides to the user, the remote and dynamic nature of a branch office (the so-called "extended enterprise") *increases* the applicability of a WLAN. As is the case with a temporary building, a branch office building or office suite tends not to be owned by the enterprise itself. Granted, this arrangement usually is in the form of a longer-term lease rather than a simple rental agreement, but the temporary nature of the relationship is fundamentally the same and there are commonly additional complications with negotiating infrastructure changes to a rented or leased facility.

Wi-Fi LANs, particularly in smaller facilities, can be remotely installed by the IS staff by providing direction to a local contractor or even an enterprise employee, which decreases or even eliminates the need for travel to remote locations. And again, when the lease term expires, the WLAN portion of the network infrastructure is portable and reusable, not buried in the walls of someone else's building.

Security Alert: The Consequences of Unmet Demand

With inexpensive, easy-to-install residential versions of Wi-Fi readily available to end users through computer retailers, consumer electronics stores, catalogers, and the Internet, Wi-Fi has been installed in many homes. This exciting market is discussed further in Chapter 10. This dynamic also has implications for the enterprise.

It's instructive to briefly review the way in which PCs entered the enterprise. Few IS staffs in the early 1980s took the initiative to deploy PCs to enterprise users. Rather, it was far more common for them to battle the proliferation of the devices until it became apparent that the fight could not be won. Enterprise IS had typically deployed a centralized and secure information infrastructure based on mainframes and minicomputers with simple "dumb" terminals deployed on the desktop. It was at the departmental and even individual level that PCs began to enter the enterprise. Users demanded the freedom and flexibility of a PC, a demand that was unmet by all but the most forward-thinking IS organizations. With PC prices falling to within the reach of departmental and individual expense budgets, it became possible to bring them into the enterprise without the involvement, or sometimes even the knowledge, of the IS organization.

The same dynamic today is playing out with Wi-Fi. Users are increasingly familiar with the benefits of Wi-Fi, often having experienced them firsthand in their homes. Whereas early PCs barely fit into departmental and individual budgets, residential versions of Wi-Fi access points can be purchased for a few hundred dollars, an amount that causes little scrutiny in most enterprises. The small size of Wi-Fi access points allows them literally to be hidden from view under a box or behind a desk, and installation is about as easy as plugging them in to an available and ubiquitous Ethernet jack.

As alluded to previously, in the absence of an enterprise IS-sanctioned Wi-Fi infrastructure, users will create their own. The problem with this grass-roots infrastructure is that individual users tend to pay little heed to the management and security requirements of the enterprise IS infrastructure. And, with an unsecured access point attached to an Ethernet port broadcasting a signal that passes easily through walls, the situation is tantamount to installing an Ethernet jack in the parking lot. Not only is the Wi-Fi network unsecure, but by attaching to an Ethernet jack that itself has no authentication mechanism, it opens access to the whole of the enterprise network, both wireless *and* wired.

This dynamic is likely to expand over time. More and more, "Wintel" laptop manufacturers—those providing devices based on Intel x86 architecture and Microsoft Windows operating systems—are providing embedded Wi-Fi adapters with their products as low-cost options. Apple Computer has been providing embedded Wi-Fi since 1999 with great acceptance. Many expect more than half of all Wintel laptops to ship with embedded Wi-Fi by the end of 2003. With departments within the enterprise rather than the central IS organization often being responsible for end-user device purchases, many choose embedded Wi-Fi in their laptops. The users of these increasingly ubiquitous Wi-Fi-enabled devices will be looking for the infrastructure needed to make this feature useful—and, as a matter of fact, most laptop vendors are happy to sell a low-cost access point with the laptop. Those who have been around networks and computers long enough recognize that the very same thing happened with the inclusion of modems and then Ethernet ports on PCs. No brand-name PC or laptop is sold, or at least used, without one to three different data access devices such as modem, Ethernet port, PCMCIA slot, and now built-in WLAN clients.

The point is that, as happened previously with PCs, the proliferation of Wi-Fi into the enterprise likely is inevitable—not from the top down, and not as an organizational initiative, but rather from the grassroots up. This occurs on a worldwide basis, from military sites to Wall Street to the smallest print shop. Wi-Fi is a disruptive technology, a revolution. IS professionals can be in the vanguard, deploying a Wi-Fi network that is as manageable and as secure as the wired LAN, or they can let the coming wave crash over them.

In Chapter 11, we detail how you can find rogue Wi-Fi networks and provide a variety of strategies for deploying a secure WLAN. While policing the enterprise and rooting out rogue networks is a prudent short-term tactic for maintaining network security, the more strategic and long-term means of addressing the current and increasing number of rogue access points is simply to preempt the incentive for individuals to deploy them by deploying instead an enterprise Wi-Fi network. After all, when was the last time someone snuck a PC into work?

Capacity Planning

Having defined a deployment strategy, the next step in the process should be to define what level of WLAN service you need to provide to the Wi-Fi users. Wireless LANs are, by their nature, a shared-medium technology. An access point establishes a coverage area or cell that provides an *aggregate* amount of throughput that is shared by all the client devices within that cell, associated to that access point. In Ethernet terms, a coverage cell is a collision domain. With Ethernet, you can define the precise number of client devices within the collision domain by choosing how many ports on an Ethernet hub will be used. With WLANs, there are, of course, no physical ports; you use the size and shape of the coverage area as a means of limiting the number of users who typically are associated to that particular access point. The means by which you can decrease (and indeed increase) the coverage size of an access point is covered in the next section on coverage planning.

With Ethernet, capacity planning is an absolute: the number of users connected to a single hub is the same as the number of users in the collision domain (assuming the hub is on its own switched segment). With Wi-Fi, on the other hand, the number of users can vary greatly as they enter and exit the coverage area. Additionally, with transmission over radio waves, throughput is subject to variation as transitory factors such as interference that decrease throughput present themselves in the coverage area. As such, capacity planning for WLANs is an approximation.

The central question that needs to be answered is: "How much throughput should, on average, be provided to each user of the Wi-Fi LAN?" Naturally, different types of users have different average throughput requirements. Warehouse and retail workers with bar code scanners have very modest throughput requirements. Office and classroom users transferring e-mail, browsing the Web, and exchanging the occasional word processor document, spreadsheet, or presentation file have greater, yet still relatively modest, throughput requirements. Finally, those transferring high-resolution graphics and layouts, CAD (computer aided drafting) files, and x-ray and other medical images have very large throughput requirements. Because the question of average-per-user throughput is essentially a division problem, one can affect either the divisor or the dividend to achieve the same quotient. The following are a few illustrations:

- **Stockroom associates with bar code scanners** For these sorts of devices, 25Kbps provides more than enough bandwidth per user. 802.11b-compliant WLANs provide approximately 5Mbps of aggregate throughput when set to an 11Mbps data rate, and provide approximately 500Kbps of throughput when set to a 1Mbps data rate. The maximum number of users per access point set to 11Mbps would be 200, with the maximum number when set to 1Mbps being 20. Few warehouses and retail locations have more than 20 associates performing bar code scans at a single time within the same collision domain. In this scenario,

the goal would be to provide physical coverage in all areas where scanning is performed with as few access points as possible—capacity is not a real issue in this scenario. As an aside, hybrid devices that serve both as bar code scanners and cordless telephones are becoming increasingly popular in these markets. The need to support voice as well as data complicates this scenario considerably and will be covered in Chapter 12.

- **Students accessing a university intranet site while in a lecture hall** While the Hypertext Transfer Protocol (HTTP) is fairly efficient, the transfer of graphics-rich web pages requires a substantial amount of bandwidth, say 300Kbps, for an acceptable user experience. This requirement becomes all the more onerous when, as part of an instructor's presentation, many students might access the WLAN at nearly the same time. With an 802.11b-compliant Wi-Fi access point providing about 5Mbps of aggregate throughput, the number of users per access point should be about 17. For a class of 85 students (not uncommon at the university level), this translates to a need for five access points in the room, which presents channel reuse challenges, as discussed in the next section. Alternatively, if the technology deployed is 802.11a or, when available, 802.11g, the aggregate throughput when set to a 54Mbps data rate will be on the order of 25Mbps, resulting in the provision of the same 300Kbps of throughput with a single access point. In this particular scenario, the very high density of users sitting in lecture hall desks renders the relatively limited range of 802.11a (and, to a lesser extent, 802.11g) a nonissue, as the single access point should be capable of covering most lecture halls, which themselves tend to be very open indoor facilities.
- **Office users transferring files** With presentation files, spreadsheet-based financial models, and even some word processing documents going well beyond 1MB in size, office workers (for whom time is, after all, money) often demand WLAN performance that compares to the switched wired connection they typically have on their desktop. For these users, their per-user throughput requirements can easily be a half a megabit per second, and ten users to an 802.11b access point may well be the right number to budget. Here again, if it's an 802.11a or 802.11g access point with approximately five times the aggregate throughput, 50 users could occupy the same coverage area and enjoy the same average per-user throughput—to the extent that 50 users could occupy the coverage area provided by shorter-range high-performance access points. Today's office cubicles are small, but not *that* small.

Naturally, the operational capacity at any given point in time is not entirely the decision of the IS department. First, defining the level of performance users can expect often requires negotiations with representatives of the user community. Remember, too, that these are the same users who have grown accustomed to a switched 100Mbps wired connection their desktop. IS professionals know that the utilization of this

connection is well less than 10 percent for the great majority of users. Nevertheless, it's not uncommon for users to want both the freedom of wireless and the level of performance that they think they need.

End-user decisions, or at least their decisions in conjunction with the decisions made by laptop vendors, also play a role in operational capacity. As discussed in the sidebar earlier in this chapter, it is becoming increasingly common for laptops to be offered with Wi-Fi radios embedded directly into the device—and for laptop purchase decisions to be made at the departmental level, not by the IS organization. The antennas for these radios are themselves embedded around the laptop's display. As discussed in Chapter 5, antennas are specially designed to transmit and receive radio energy within a certain frequency band. With design cycles for laptops of approximately a year and a half, most antennas embedded in today's laptops are tuned to 2.4GHz, the frequency band of 11Mbps 802.11b, not the 5GHz band of 54Mbps 802.11a. The irony is that the laptops with wireless embedded that are demanded by users and departments complicate the ability of the IS department to provide them with the performance that they think they need.

It is, by the way, typical for 802.11b Wi-Fi networks to be deployed for 11Mbps coverage areas. Given their relatively short range of 802.11a at their maximum data rate of 54Mbps, it is more typical to plan for one of the lower supported data rates that provides for greater range—although the newness of the technology makes generalizations like this difficult.

The principal way to increase per-user throughput is to decrease the number of users contending for the aggregate throughput provided by the access point. This limiting of users is typically accomplished by decreasing the size of the coverage cell. Two major implications arise from this.

The first is that it doesn't come free. The obvious implication of decreasing the coverage area of an access point is that more access points are required to cover the same given physical area. Doubling the amount of throughput provisioned for each user doubles the cost of the access points and the deployment thereof.

The second implication is that deploying for higher per-user throughput can simplify deployments—or complicate them. As discussed in Chapter 1, the legacy environments for WLANs are similar to those described in the earlier stockroom example—a relatively low density of users with low bandwidth requirements. Accordingly, the physical planning for WLANs focuses on achieving coverage in all required areas with the fewest number of access points possible. After all, as recently as 1999, access points cost more than $2000 each, while leading performance devices today are approximately one third of that price. Deploying for high per-user throughput eliminates the need to optimize access point range.

You must also consider that the cost of deploying an AP includes the labor cost and, commonly, the cost of deploying additional Ethernet cable and access to AC power (although APs exist that do not require separate lines for power and data). High transmit power, receive sensitivity, and antenna gain are unnecessary when limiting the number of users in the collision domain through *decreased* cell size.

A high density of access points does, however, present other problems. Not all vendors provide features like transmit power control settings that are designed to decrease coverage area. Antenna attenuators that decrease the gain of an antenna, and therefore cell size, can be expensive and are only a possibility when using antennas with connectors. Finally, when spacing access points close together, channel reuse problems become more acute, particularly in the narrow 2.4GHz band that allows for just three nonoverlapping channels.

Coverage Planning: The Site Survey

In the first chapter of *Baby and Child Care,* his seminal book on child rearing, Dr. Benjamin Spock famously started with, "Trust yourself, you know more than you think you do."

The idea was to reassure concerned and even frightened first-time parents that they should trust their intuition when raising their children. After all, parents had managed to raise their children before instruction manuals and trained professionals existed. Parents then and now draw upon their experiences, their intuition, and the advice of other parents. With all that said, in some more challenging situations, parents look to professionals for guidance and assistance. And indeed, as Dr. Spock has told you, they buy and read books on the subject.

Today, faced with deploying a WLAN, IS professionals in the enterprise are a little like first-time parents, competent and effective people confronted with a new and unfamiliar challenge. IS professionals are typically well versed in *wired* network architectures, the tools designed for managing the *wired* LAN, security policies that presume physical ports, and, of course, even the bend radii for various types of fiber-optic and coaxial cable. All of which, at least at the surface, have little to do with Wi-Fi.

On the other hand, IS professionals, like anyone else living in an industrialized country in the twenty-first century, have a lifetime of experience with radio waves. We watch television and listen to AM/FM radio. We might have even had a CB radio (although typically we choose not to admit it). Walkie-talkies, pagers, cell phones, baby monitors . . . we've grown up with radio and we live with it still. We know intuitively that radio waves go through walls but that the signal is weakened when they do; that subtle movements and changes in position can have a huge impact upon how well a signal is received; that a signal gets weaker as it gets further from its transmitter; that when two signals are at similar frequencies, they can interfere with each other; and that, just as visible waves of light can be blocked, creating shadows, radio waves can be blocked, causing a signal to disappear as we drive through a tunnel. Trust yourself, you know more than you think you do. This is not to say that an IS professional should lumber into a Wi-Fi deployment unaware, or that there's no difference between a wired and a WLAN. Rather, it's meant to point out that performing the tasks that are specific to a Wi-Fi network, and that are necessary for a successful deployment, can all be learned and that you probably have a bigger head start than you think.

If the goal of capacity planning is to provide users with what they need, the goal of coverage planning is to provide them with what they need *where* they need it. This relates back to the various deployment strategies—some areas will be designated for

WLAN deployment and others will not. Coverage planning is often referred to as a *site survey*, a process whereby an individual or group gathers data and then makes specific recommendations as to the types of access points, antennas, and other equipment to be installed and the specific locations for these installations.

A site survey takes into account the design of the building and its construction materials (ascertained through blueprints and floor plans as well as direct examination), the traffic patterns within the facility, the sorts of barriers likely to be encountered in the facility, the range and coverage pattern capabilities of the access points to be used and the flexibility of those capabilities, the technologies (802.11b, 802.11a, or both) and resulting throughput channels available to them, and, of course, the capacity plan.

When Is a Site Survey *Not* Necessary?

Before answering the question, first let's better define the term *site survey.* In the most elementary sense, a site survey is a simple look around a facility before placing an access point. In a home, it can be a matter of choosing on which bookshelf to place the access point. In the more extreme cases, a site survey can take days and require you to hire experienced and trained professionals who tend not to work cheaply, and who provide the network administrator with a large binder full of information about WLAN network element placement. For the purposes of this question, we define a site survey as requiring the services of someone specially trained in doing them, which tends to mean hiring a consultant or a reseller of WLAN and other network hardware.

Frankly, those with long-time experience with WLANs sometimes tend to overemphasize the need for a site survey, probably more so in SOHO environments than in a larger enterprise network where network unreliability readily converts to operational inefficiencies and lost profits and revenues. The tendency of some WLAN professionals to over-optimize a site survey may result from their experience in challenging applications like retail locations, warehouses, and hospitals. Naturally, these are some of the earliest adopters of WLANs and the types of installations industry veterans have a great deal of experience with—to the exclusion of more recent enterprise adopters that are generally less challenging environments for WLANs. Another possibility is that there is a fairly lucrative market for the professional services needed for what are sometimes unnecessary site surveys.

As a general rule, if a single access point can cover a facility and provide the per- user throughput required, and the facility has no server, a site survey is unnecessary. Remember that even 802.11a access points, which provide less coverage than their 802.11b counterparts, provide an approximate 50-foot coverage radius at their 54Mbps maximum data rate. With an omnidirectional antenna providing a circular coverage pattern, the resulting coverage area is 11,000 square feet. With an 802.11b access point with a 100-foot coverage radius, the area grows to more than 30,000 square feet.

The point is that even 11,000 square feet is larger than most homes, most small offices, and most branch offices. Assuming these facilities are made with standard building materials like wood, drywall, and plaster for interior walls and don't have an inordinate number of interior fixtures like file cabinets and whiteboards that are unfriendly to radio waves, they can typically be covered by a single access point. Even in situations where more than one access point is needed either for capacity purposes or to cover the corners or recesses of the facility, their placement is fairly intuitive. Remember that even the 2.4GHz band provides for three channels. You can place as many as three access points in a facility (it's good practice to keep them at least ten feet apart to avoid interference) without any concern for interference between the devices.

In short, even in situations where up to three access points are required to cover the facility, you may well be able to dispense with a formal site survey. This isn't to say that you should indiscriminately install access points without thinking; it means that with a little planning, study, and common sense, you can successfully deploy your Wi-Fi LAN.

Internal and External Building Design

"They just don't build 'em the way they used to" is a common refrain heard regarding buildings. Hallmarks of buildings from the first part of the twentieth century and before are brick or even stone external walls, plaster and lath internal walls stretching from floor to ceiling, and high plaster ceilings. In North America at least, buildings from the postwar era are a very different story. Exterior walls generally are thinner, predominant construction material for interior walls is drywall, open spaces separated by cubicles are more common, and larger windows and suspended ceilings are the norm.

Although people can and do decry the perceived decline in building quality, the newer buildings are a lot more friendly to Wi-Fi installation. In general, the more dense the construction material, the more it prevents RF energy from passing through it. This matter of energy loss is referred to as *attenuation.* Wood, drywall, cubicle walls, room partitions, and the like have a relatively high amount of air in them, whereas brick, cement, stone, and thick plaster walls have less air in them, and also tend to be thicker. Metal, such as the exterior metal walls of a warehouse or hanger, or even the metal studs used today for interior walls instead of wood, presents a special problem because it not only stops a signal, but reflects it, creating the multipath propagation discussed in Chapter 5.

Understanding the effect various building materials have on radio energy makes for a good starting point when surveying the facility to be covered. Through either blueprints or, better still, direct physical inspection, you should familiarize yourself with the types of construction materials found in the facility.

- Avoid planning for penetration through exterior walls, as they typically degrade a signal to a large and unpredictable degree. This makes the resulting exterior coverage area variable in performance and reliability. If an exterior coverage area is desired (as is often the case with university and corporate campuses), antennas should be installed outdoors specifically for these coverage areas.
- Plan for little attenuation when installing an access point in an open office environment, like the types populated with cubicles.
- When installing a Wi-Fi LAN operating in the 2.4GHz band, plan for penetration through most interior walls, including those made from drywall, plaster, and even cinderblocks, although they provide increasing levels of attenuation. The metal studs often found in interior drywall in commercial buildings can introduce a level of unpredictability and multipath when at a high angle of incidence to the transmitter. The 5GHz waveform of 802.11a Wi-Fi LANs is absorbed and distorted by common materials to a greater degree than is the 2.4GHz waveform, due to the differing length of the waveforms. With a physical length of around two inches, the 5GHz wave is about half the length of the 2.4GHz wave, causing it to deteriorate more extensively as a function of time and as a function of coming into contact with structural elements. Generally, you can plan for penetration through drywall and plaster but typically not through cinderblock when working in the 5GHz band. Note, however, that the relative newness of 802.11a results in a much smaller body of empirical installation data in the world of WLAN.
- While you can plan for coverage through walls that are *partly* made of metal, you cannot assume penetration through all-metal walls. Indeed, due to the multipath-inducing properties of metal, you should plan around them.

The operative words here are "plan" and "assume." That is, this is just the first step in the process, which is then followed by an actual physical verification of these assumptions. The documentation of these assumptions is usually a site plan in which provisional access point placements are made to a copy of the building's floor plan. This is a very useful and arguably indispensable tool for implementations. It is also useful for establishing a budget, as you can get a fairly accurate estimate of the number of access points, antennas, cables, and other accessories needed at this stage. A compass, set to the correct scale of the floor plan, is ideal for estimating omnidirectional and hemispherical coverage areas.

As shown in Figure 9-1, in the 2.4GHz band, three nonoverlapping channels are available for 802.11b and, when available, 802.11g. If a facility can be covered with three or fewer access points, co-channel interference is not a problem, which significantly simplifies a deployment. A single access point, and certainly three, will cover contemporary office floors, which tend to have only partial cubicle walls rather than floor-to-ceiling walls, which tend to attenuate RF energy and reduce coverage.

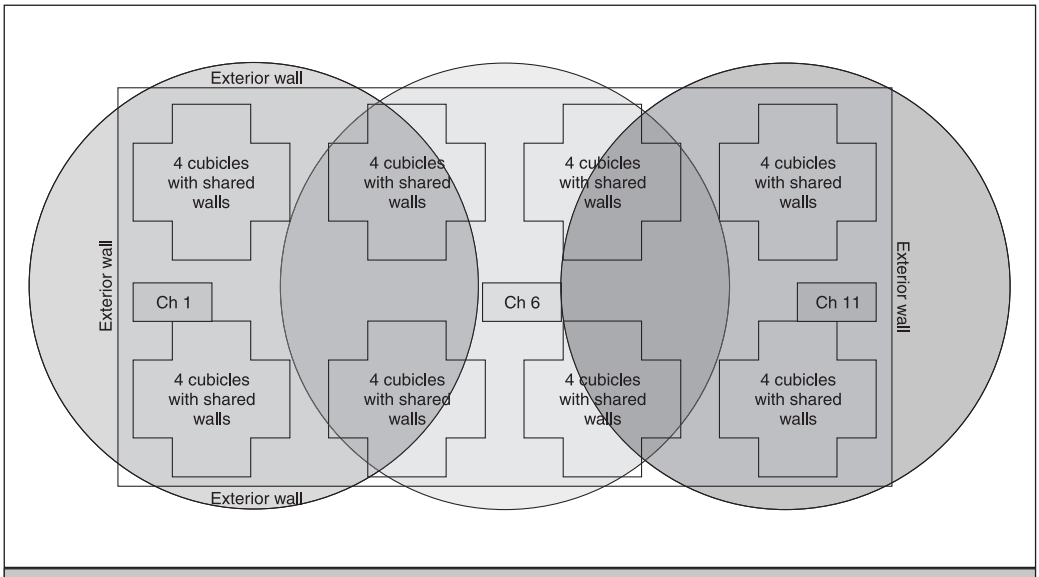

Figure 9-1. Top-down view of an open office floor

Retail stores and warehouses are among the more challenging environments for RF coverage, as shown in Figure 9-2. Depending upon elevation, material selection, and angle relative to the transmitter, shelved merchandise can very effectively block RF energy from one row to the next. The metal shelves themselves reflect RF energy and create multipath propagation, which drives down performance. An effective strategy for deployment in these facilities is to use patch or Yagi antennas to direct the RF energy in a tight beam down the rows, thereby covering the areas associated with inventory control and minimizing the reflections from the metal shelves.

The propagation characteristics of the 2.4GHz waveform allow for penetration through walls made of many types of building materials—even the cinderblock walls often found in primary and university classrooms. An 802.11b and likely an 802.11g access point with an omnidirectional antenna placed against a wall can often cover two classrooms. Given the high user density consistent with classroom environments, the per-user throughput provided by this deployment could be quite low. With only three nonoverlapping channels available, the isolation of cells set to the same channel becomes an issue. Note that the coverage area provided down the corridor overlaps slightly with a classroom coverage area set to the same channel. In more densely deployed environments like this one, minimizing co-channel interference rather than eliminating it entirely can be the goal.

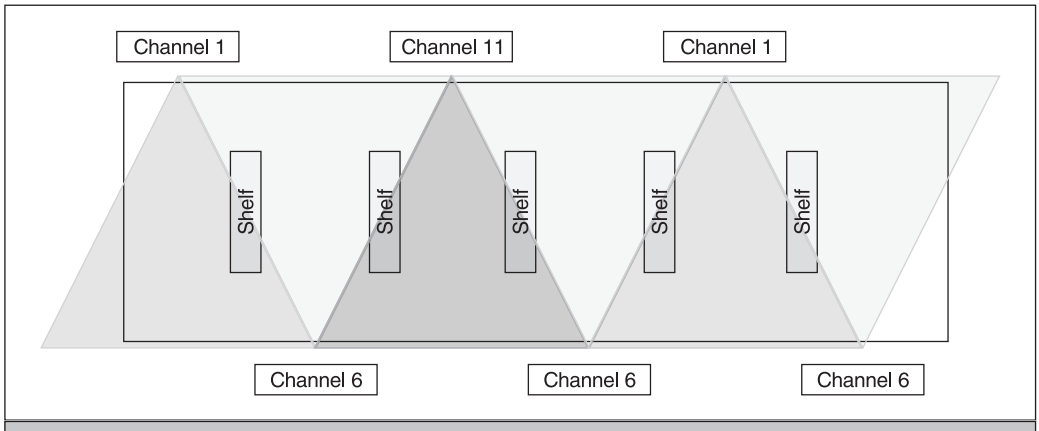

Figure 9-2. Top-down look at a warehouse floor or retail location with shelving units

The 5GHz waveform is attenuated by common building materials to a greater degree than the 2.4GHz wave. Whereas a 802.11b access point can be installed such that it covers more than a single room, the coverage of an 802.11a access point more typically is limited to a single classroom—particularly when constructed of materials like brick or cinderblock. Given the high user density consistent with classrooms, this characteristic can actually be beneficial in that it helps to limit the cell size and thereby provides for a higher level of per-user throughput. Note that the eight channels available in the UNII-1 and UNII-2 bands combined decrease, if not eliminate, channel reuse concerns. Even with a very large number of access points deployed, the large number of channels (coupled with the more limited cell size) allows for a deployment with no overlapping cells set to the same channel, thereby eliminating any performance-degrading co-channel interference.

Some objects and building materials are essentially impenetrable to radio waves of any frequency, as shown in Figure 9-3. Elevator shafts with a large amount of steel and, to an even greater degree, the x-ray rooms with lead-lined walls that are commonly found in hospitals are best planned around. 802.11b access points with omnidirectional antennas can cover a number of examination rooms while this same type of access point, installed in the x-ray room, covers just that room. Patch antennas with a wide-angle coverage pattern can be used to fill in areas not covered by omnidirectional antennas. Note that while the three available channels in the 2.4GHz band are a limiting factor, access points can be deployed such that full coverage can be achieved without any overlapping cells set to the same channel.

Multistory structures like office towers, hospitals, and university classroom buildings introduce a third dimension to coverage planning. The 2.4GHz waveform of 802.11b

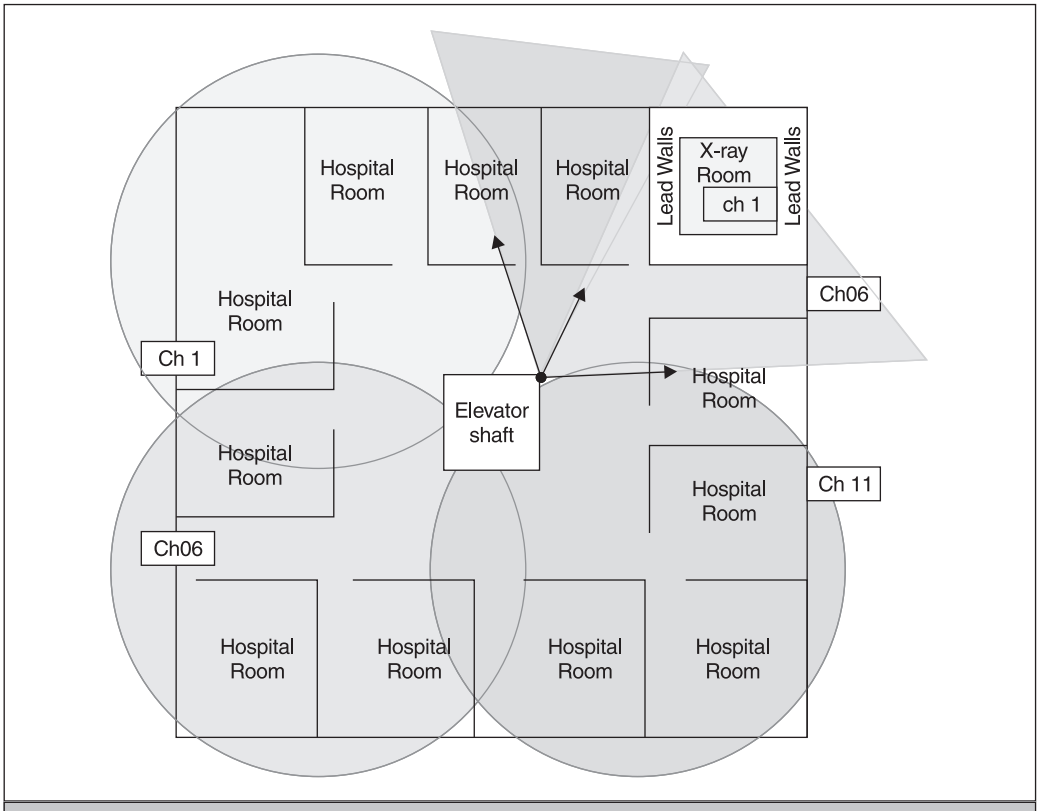

Figure 9-3. Objects and building materials that are impenetrable to radio waves of any frequency

and, when available, 802.11g will pass through both walls and floors. The 5GHz waveform of 802.11a will pass through both as well, but to a lesser degree. For 2.4GHz Wi-Fi LANs in particular, you must avoid overlapping cells on the same floor and on adjacent floors. Even with only three channels, this can be achieved through careful three-dimensional planning.

Placement Options

Choosing the best locations for access point and antenna placement involves a number of different and sometimes competing factors. Optimal locations from a propagation perspective may be aesthetically or economically unacceptable. Budgetary constraints may result in access points with suboptimal range and reduced antenna options. Every building is going to present different parameters suggesting different placements, but some general rules do apply.

Roaming

Strictly defined within the Wi-Fi standards, *roaming* is the process whereby a client can move from the coverage area of one access point to the coverage area of the next access point. As discussed in Chapter 5, this is accomplished through a client-side process of scanning for available access points and, if better performance can be had when associated to another access point, disassociation followed by association with the other access point. Roaming can, however, also be thought of as the means by which you can scale a Wi-Fi deployment to cover an area of almost any size. Indeed, as will be discussed in Chapter 14, Wi-Fi deployments covering entire towns and neighborhoods are already in place, with citywide deployments within the realm of technical possibility.

You should consider the infrastructure to be a web of interconnected access points that, in the aggregate, provide complete, uninterrupted coverage. Roaming is the capability that allows a client to "view" access points in the same manner. That said, the process of roaming is neither perfect nor instantaneous. Frequent roams increase the possibility of noticeable performance degradation, particularly when running latency- sensitive traffic like voice and video. And, of course, each "coverage area" is actually an access point that costs money. Given this, you should minimize the number of roams necessary to cover the required area by maximizing the cell size of each access point—within the context of the capacity plan.

Ceiling placements tend to work best. By placing access points or antennas on or above ceilings, circular cells that maximize the coverage area of the access point can be set up with omnidirectional antennas, the most common type available. Placement on the ceiling gets access points and antennas away from people, minimizing intentional or unintentional contact. Some access points can be hidden above suspended ceilings (see the sidebar on plenum considerations) with only antennas visible, a positive feature from both aesthetic and theft-deterrence points of view. Even access points with nonremovable antennas and plastic cases can sometimes be mounted on ceilings.

When working with access points with connectors that support auxiliary antennas, you have some flexibility in terms of access point placement. Access points can be remotely located in places like wiring closets that provide for centralized management and theft protection. On the other hand, as discussed in Chapter 5, substantial loss in gain results from cable runs. This cable loss can negate any gain provided by the antenna. In facilities with suspended ceilings, it's more typical to place an access point designed for these locations (see "What Is the Plenum?" sidebar) near the antenna than it is to suffer the cost and cable loss associated with remote access point placement.

Ceiling placement allows the RF energy to radiate down to the floor below and, in some cases, can even provide a "bonus" coverage area on the floor above. Access points may be placed in ceiling centers or, if more than a single access point is necessary for coverage or capacity, spaced at intervals that in the aggregate provide for full coverage (refer to Figure 9-1).

Mounting access points on desktops or, better still, on the top rails of cubicles provides benefits similar to ceiling mounts. This sort of installation is common when working with lower-cost access points that are ill-suited to ceiling installation or when an installation is likely to be temporary.

What Is the Plenum?

The National Electric Code (NEC) defined in 1999 the *plenum* as being "…a compartment or chamber to which one or more air ducts are connected and that forms part of the air distribution system." In other words, the plenum is typically the space above a suspended ceiling where things like heating and air conditioning duct work runs and overhead lights are installed, an area that's ideal in many ways for the placement of Wi-Fi access points.

Because of the types of things installed in the plenum, it's an area that is subject to fire codes. In the event of a fire, this area is ideal for the spread of flames and, more to the point, noxious fumes and poisonous gasses, which are a far greater cause of fatalities in fires than flames themselves. The types of materials you can put in the plenum are therefore restricted.

Because the plenum is subject to *local* fire and building codes, there is no one national or international standard for what types and amounts of materials are acceptable for placement in the plenum. Indeed, some municipalities define the plenum to include not only the space above a suspended ceiling, but also an area extending some number of inches *below* the suspended ceiling. Still, you can make choices that maximize the likelihood of steering clear of problems with the local building inspector or, much worse, creating a hazard in your workplace.

Although there is no universal standard for plenum rating, there is a very good substitute. Underwriters Laboratories has developed a standard, UL2043, titled "Fire Test for Heat and Visible Smoke Release for Discrete Products and Their Accessories Installed in Air-Handling Spaces." This test, and resulting compliance to the standard, does not address the toxicity of the fumes released when the device is burned; rather, it addresses the rate at which they burn and the quantity of energy released and then ignited. Still, if a product is certified to UL2043, it's a good bet it'll meet with local building codes. Cisco Systems, www.cisco.com, is alone in certifying selected access points to this standard.

Absent UL2043 certification, the rule of thumb is to avoid access points with plastic cases for placement in the plenum. Anyone who's burned a model airplane or unfortunate army man knows that plastic burns and, when it does, releases dense black smoke and a variety of harmful gasses. On the other hand, if a device has a metal case, it'll most likely be acceptable for placement in the plenum. Even devices with a small amount of plastic, such as connectors, labels, and "rubber" feet, tend to be acceptable. In addition to Cisco, other vendors like Enterasys, Proxim, and

Symbol Technologies provide metal-cased access points or access points with plastic cases that can be removed to reveal an inner metal case.

In short, when installing an access point above a suspended ceiling, make sure it has either UL2043 certification or a metal case.

In buildings where ceiling mounting is impractical, would represent an unacceptable disruption to normal operations, or is considered to be aesthetically unpleasing, wall mounting is an increasingly popular option. With omnidirectional antennas installed, 2.4GHz Wi-Fi access points mounted on walls can often cover two rooms. For 5GHz Wi-Fi devices, the attenuation associated with the waveform through walls eliminates the two-room option. For either 2.4- or 5GHz Wi-Fi access points, patch antennas can be used that direct the RF energy from the wall across the room. By placing multiple access points on walls, complete coverage from wall-mounted units can be achieved in all but the largest rooms.

The Physical Site Survey

With a capacity and coverage plan complete, you can test your assumptions. Now is the time to actually place access points and selected antennas in their provisional locations and test for coverage. In the same way that an actual product can be very different from its data sheet, a building can be very different from its floor plan. And signal propagation in practice can be inexplicably different than it is in theory—this is, after all, radio. So, before buying equipment and permanently installing it, it's very wise if not mandatory to do test installations at most, if not all, of the provisional locations defined during the coverage planning process.

NOTE A key thought to remember is that not even the most experienced RF engineers will trust their eyes as to what *should apparently* seem to work in the realm of RF propagation. While you can have indicators and even a reasonably developed instinct for how a radio will perform in a certain location, the longer you are in this industry, the more careful you tend to become about acting on the assumption "I can't imagine it wouldn't work just fine in this room." A second interesting thought is that the trend now for high-value networks is to occasionally repeat the site survey, because the general tendency for networks that use 2.4GHz, and many other frequencies, is degradation of the radiating environment over time due to co-channel interference (someone else in the vicinity also broadcasting in the same frequencies), adjacent channel interference, or the alteration of the physical environment. Perhaps one of the best examples of this is in the financial markets in New York City, where site surveys are completed *on a weekly basis* just to ensure one other very important item—the absence of rogue equipment.

The frequency and scale of your site surveys dictate to a degree how comprehensive your site survey toolkit should be. There are, however, some basic tools that you should have that will make this process easier and more effective.

First, at minimum, is a vendor-provided site survey tool. Most Wi-Fi vendors provide site survey tools of varying capabilities with their client adapters utilities. As shown in Figure 9-4, a site survey utility reports the access point to which the client adapter is associated. It reports the strength of the signal and the resulting data rate supported as well as the ambient noise level. Some site surveys incorporate what amounts to a ping test to measure the number of IP packets lost during a transfer. You'll want to ensure through your site survey efforts that you have not only good signal *strength*, but also good signal *quality*. This is important for a number of reasons, not the least of which is that you can consider installing an access point at a location where a good amount of RF energy in the correct frequencies is received but, for various reasons, will not carry an appropriate ratio of recognizable bits to degraded bits.

As part of your capacity plan, you will have established the required data rate to be provided by the access points and the location and number of users for a given access point. With a site survey tool, you can ascertain not only the associated data rate but also the reliability of that data rate by taking into account more qualitative data. like

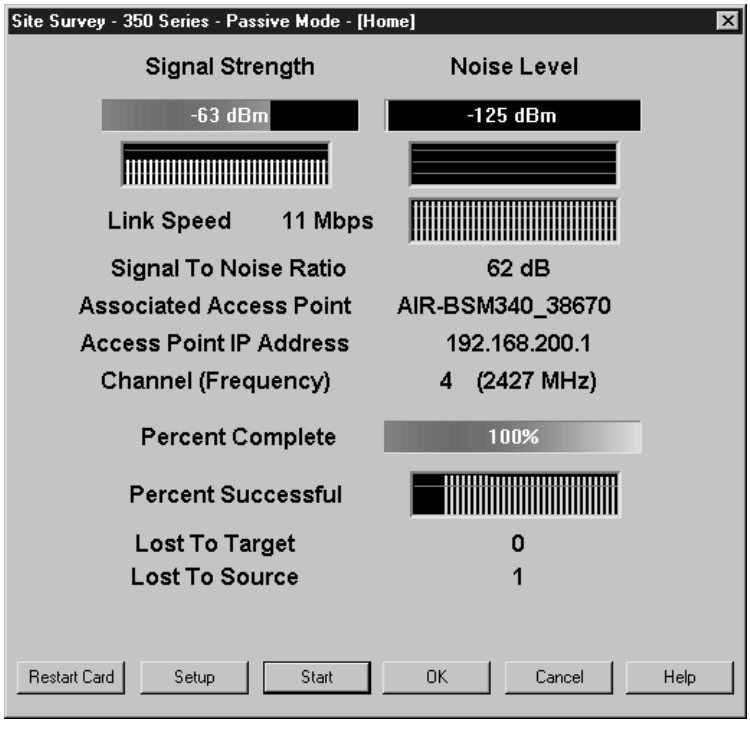

Figure 9-4. Site survey utility report

signal strength and packet loss. As shown in Figure 9-5, some site survey tools report a subjective level of the signal quality, such as Excellent, Good, Fair, and Poor, either in lieu of, or as an option to, the decibel scale, although most experienced WLAN professionals prefer that the tool simply report back the raw data. As discussed in Chapter 5, the signal-to-noise ratio (SNR) is a useful metric for assessing the reliability of the link at a given data rate. As an example, at 11Mbps, it is typical to provide for a link margin of at least 10dBm.

For the Serious Site Survey: AirMagnet

Most Wi-Fi client adapters come complete with a site survey tool that reports basic information, such as the data rate, signal strength and quality of the associated access point, the ambient noise floor, and the resulting SNR. Some provide other rudimentary but useful tools. For many site surveys, these tools may be all that's needed for a successful deployment. And since they come free with a client adapter, the price is certainly right.

More advanced tools may be in order for more involved site surveys, or for those who do a lot of site surveys or need to integrate their 802.11 equipment at the highest levels of network integration. In the freeware category, NetStumbler, www.netstumbler .com, is designed as both IS professionals a site survey tool and, as per the NetStumbler web site, "overly curious bystanders" and "drive-by snoopers" as a means to "pick up ladies." The site is worth checking out if for no reason other than entertainment value. The principal function of NetStumbler is to search the airwaves for access point beacons and then display them. In the hands of a hacker or someone looking for free access to the Internet, it can be dangerous. On the other hand, it can be useful to the IS professional to check for unprotected and rogue access points (more on security in the next chapter) or as a tool for checking multiple AP coverage. NetStumbler supports a variety of client adapters and runs on most Windows desktop operating systems and Windows CE.

The most full-featured site survey tool is AirMagnet, www.airmagnet.com, from a company of the same name. AirMagnet runs exclusively on Windows CE, meaning that it's designed for operation on a PDA such as a Compaq iPaq. Although PDAs make ideal devices for site surveys due to their small size, their 16-bit PCMCIA interface (rather than 32-bit CardBus) limits current support to 802.11b and not 802.11a client adapters. To describe AirMagnet as a site survey tool is a bit of an injustice, because the full scope of this tool includes a variety of capabilities for security and performance monitoring—indeed, the product has been described as a "Swiss army knife for wireless LANs." For site surveys, it provides detailed information on the whole RF environment plus packet-level data in a helpful graphical format. All this functionality comes at a price, $2495 to be exact. Still, professionals require professional tools, and for the right people, AirMagnet could be the right product.

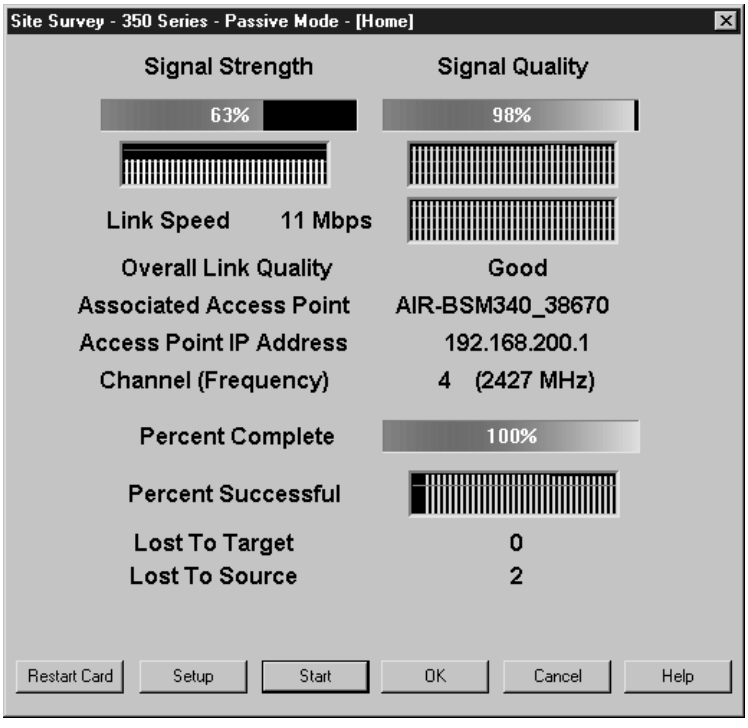

Figure 9-5. Site survey with subjective levels of quality

In addition to site survey software, other tools are necessary for a site survey. Since you'll be placing access points or antennas in temporary positions and may need to make a few adjustments until the optimum location is found, you need a quick yet sturdy means of mounting the devices to a variety of surfaces that also enables you to remove them without damaging the surface they are attached to. For some surfaces, duct tape (naturally) works, although it has the tendency to remove paint. An increasingly popular alternative to duct tape is to use zip ties, which can be used in areas where the access point will reside upon fairly new paint, or on top of a wallpapered surface.

Once you have identified access point installation locations, you need some means of marking these locations. Firms specializing in site surveys sometimes use small flags with their company logo on them—they're effective for marking locations as well as advertising their company. Brightly colored tape works just as well but without the commercial benefits. Finally, you need a tape measure or, better still, a measuring wheel for recording distances.

An important item to remember when performing a larger site survey as well as smaller site surveys is to record the location of the APs in some manner that will make sense a year later both to you and, even more importantly, to other individuals.

Documentation of where APs are located is a very important item because you'll need that documentation for future reference as you build out your 802.11 network, as well as for troubleshooting and security purposes. One of the most indispensable tools for a site survey is a small camera, which is used to record the actual AP along with some wider angle shots of where it is deployed. Obviously, a digital camera is an even better asset than a camera that uses film. Veterans of site surveys will tell you that the first indicator of a professional installation (after how well the 802.11 elements perform) is how well the deployment is documented.

You have to take into account the differing capabilities of the client devices that will be using the wireless infrastructure once it is deployed. Client adapters from different manufacturers can vary by as much as 5dBm in transmit power. Different client-side form factors allow for different antenna gain, which adds further variance to client capabilities as a whole. In areas that likely will include different types of clients, you're well advised to plan for the lowest common denominator, the client types with the lowest capabilities.

Naturally, if the access points you'll be working with support external antennas, you should have a variety of antennas with you to try out allowing you to select the one that provides the best coverage for that particular installation location. In general, you'll want to have at least two omnidirectional and two patch antennas, each with differing amounts of gain in order to determine the optimal AP configuration. Variable or rotary attenuators are small devices that connect between the antenna and the access point and allow you to reduce (typically in 1dBi increments) the gain of a given antenna. This allows you to carry a smaller number of antennas and to approximate the range of lower-gain antennas. While this is a convenience, it doesn't take into account the larger beam width associated with lower-gain antennas, which could result in some unexpected results when the actual antennas are installed. In short, attenuators can be useful but are not necessary for most site surveys.

For the convenience of both the site surveyor as well as those in the facility when the site survey is being conducted, a battery pack is a useful addition to the site surveyor's toolkit. Temporarily running access points off of DC battery power saves you from long extension cord runs from AC outlets to the provisional access point locations. Extension cords running along floors and up to ceilings are unsightly, inconvenient, and even present a safety hazard, and in some cases, the site survey is performed before there is local AC power installed, such as in new buildings or highly renovated floors within a building. The complication is that battery packs specifically designed for site survey use with access points tend not to be readily available off the shelf. Added to this is the reality that differing access points run off of different voltages, ranging from a low of 12VDC to as much as 48VDC. One source for access point battery packs as well as other site survey tools is TerraWave Solutions, at www.terra-wav e.com.

WI-FI MANAGEMENT IN THE ENTERPRISE

After you have deployed a Wi-Fi LAN, it needs to be managed just as a wired LAN has to be managed. It is in this area that significant differences in capabilities between low-cost access points and those designed for the enterprise become most apparent. The

management capabilities of lower-cost offerings assume the deployment of a relatively small number of devices and, accordingly, provide little in the way of large-scale manageability. The management interface is typically limited to the browser, with no command-line interface (CLI) available—a very real complication if you prefer to write scripts to automate tasks for a large number of devices. Similarly, lower-cost devices often don't provide the Simple Network Management Protocol (SNMP) support required for operation with network management software (NMS) like Hewlett-Packard's OpenView and Computer Associates' Unicenter that is used with wired networks.

Maintaining the Infrastructure

Given the relative newness of, and the rapid rate of change in, the Wi-Fi industry, it's common for vendors to provide software upgrades as frequently as on a quarterly basis. These software upgrades can provide bug fixes, new features, refinements to existing features, and the standardized implementation of features previously released as proprietary offerings. In short, it's typically in the user's best interests to increase the functionality of their Wi-Fi infrastructure by upgrading the firmware. This means that you'll be performing software upgrades on a frequent basis across the whole population of deployed access points.

Access points developed specifically for enterprise deployments provide the management features necessary for the deployment of large numbers of devices. Like wired enterprise switches and hubs, they typically provide not only a browser interface but also a CLI, which enterprise IS professionals often find to be more efficient and more conducive to scripting, partly because it enables them to cut and paste commands from one element to another, or an array of others. The browser interface on enterprise access points tends to provide far more features designed for enterprise scalability than devices designed for SOHO deployments. While the vendors of lower-cost access points focus their efforts on ease of use, enterprise access point vendors offer features like the ability to replicate a single access point configuration across all other access points, the ability to download a firmware upgrade or configuration file from a centralized server that the access point is directed to by a link provided by a BOOT-P or DHCP server, and the ability to rekey passwords and SSID information.

In addition to the CLI and browser interface, enterprise access points provide a similar and increasing level of SNMP support that IS professionals would expect from wired LAN infrastructure devices. The management information bases (MIBs) provided by enterprise access points can be compiled by the same NMS that enterprises use to manage the wired LAN. This allows IS departments to leverage both their monetary and training investment in existing NMSs as well manage the wired and wireless network as it should be—a seamless and cohesive unit.

Regrettably, little of the remote management capabilities found with access points are found with client adapters, even client adapters provided by vendors focusing on the enterprise. Although emerging enterprise architectures are designed to require less frequent client-side configuration changes and firmware upgrades, the need to perform

these tasks is both very real and will occur on at least an annual basis. The task is further complicated by the need for user involvement, remote locations, and the greater number of clients relative to access points. Although some solutions are available today (see the sidebar "Mobile Manager from Wavelink"), there are few client-side management capabilities provided by hardware vendors and third parties to address the need for client-side management.

Mobile Manager from Wavelink

Although you can manage your WLAN using traditional network management software like HP's OpenView or CA's Unicenter, there is a product available specifically designed to manage WLANs. Mobile Manager from Wavelink Corporation, www.wavelink.com, supports access points from leading enterprise vendors like 3Com, Cisco, Intel, Proxim, and Symbol Technologies.

Mobile Manager, and the more feature-rich Mobile Manager Enterprise, can be used to automatically detect the installation of new access points on the network using SNMP. Once detected, you can send predefined configurations to all the access points, which saves you from having to manually configure each device one at a time. After you install Mobile Manager, it can be used to monitor the WLAN, providing individual access point utilization data, failure alerts, and a general log file. Mobile Manager can be configured to send alerts via pager and also to other NMSs, better integrating the WLAN with the wired LAN. While these sorts of capabilities are available from vendors like Cisco and Symbol (indeed, Wavelink collaborated with Symbol to develop its WMNS management system), Wavelink provides the ability to manage an access point environment with products from two or more companies.

Although management software like Wavelink is very useful in configuring not only a large number of APs but also APs sourced from multiple vendors network administrators and purchasers of 802.11 equipment should keep in mind that not all APs have the same level of sophistication, ease of management, reliability, and security. The 802.11 standard ensures that the radio itself, and the MAC layer will conform a minimal level of functionality and interoperability but it by no means ensures that all 802.11 equipment is created the same.

Mobile Manager works in conjunction with Wavelink's Avalanche, a product that provides similar management capabilities to client devices. With Avalanche, firmware upgrades and configuration changes can be made on a global basis across the RF and to the individual clients without any end-user intervention. However, Avalanche is expensive, putting this very useful capability out of the reach of even many enterprise organizations.

For enterprise deployments, Wavelink and other wireless-specific NMSs that may come to market are worth consideration.

Monitoring the Infrastructure

As is the case with maintenance, there are considerable differences between lower-cost devices intended primarily for SOHO use and 802.11 devices intended for enterprise deployments. Enterprise access points provide detailed association lists and logs so that IS professionals can be continually aware of the user and bandwidth load on each device; indeed, one of the key differentiators between 802.11 access points used in larger enterprise organizations compared to those used in small deployments is the amount of resolution a network administrator has available for these devices. In the same way that the site survey verifies and helps to refine the coverage plan, ongoing monitoring of access point traffic allows you to verify and refine the capacity plan. On enterprise access points, this status data is available through the CLI, the browser, or, via a compiled MIB through general and wireless-specific NMSs.

CHECKLIST

In this chapter, we focused specifically on the unique nature of Wi-Fi deployment in the enterprise. We defined the enterprise and outlined the various Wi-Fi deployment strategies that enterprises are using. We discussed capacity planning, coverage planning, and the need for a site survey. Finally, the chapter covered the types of management capabilities provided by differing access points. Some key points follow:

- ☐ For the purposes of this chapter, the enterprise is defined as being not only large commercial entities but also relatively large educational and governmental institutions—essentially, any organization of a scale that requires a dedicated IS organization.
- ☐ There are various ways that enterprises go about deploying WLANs. Some deploy Wi-Fi only in the areas where it is perceived to be needed most, like conference rooms and other public spaces. Campus-based organizations, such as universities, may deploy one complete floor or building at a time. For some enterprises, initial Wi-Fi deployments are used only for temporary workgroups or for buildings that are to be occupied on a short-term basis. Finally, geographically distributed enterprises will first deploy Wi-Fi in remote locations, like branch offices.
- ☐ A Wi-Fi LAN is, by its very nature, a shared-medium technology, where all users associated to an access point share the aggregate throughput provided by that device. Capacity planning entails planning for the maximum number of users per cell to provide on average a reasonably predictable per-user throughput. Different applications, from retail to classroom to office deployments, tend to call for differing levels of per-user throughput. Since collision domains are restricted by limiting the coverage area provided by an access point, any capacity planning is more of an approximation than an absolute.

- The very different nature of facilities calls for different coverage plans. Various building materials, floor plans, and internal structures have effects on RF propagation characteristics, which in turn are different for the 2.4- and 5GHz waveforms. There are various coverage strategies associated with different types of facilities, including offices, classrooms, and warehouses.
- Different placement options for access points are available. Ceiling mounts are the most common and generally provide for the largest coverage area, with a circular, omnidirectional coverage pattern. Ceiling placement can also result in coverage in the floor above, which may be intentional or unintentional, but in either event must be measured and accounted for. Wall mounting is an increasingly popular option since installation tends to be less disruptive to operations and less aesthetically obtrusive.
- The site survey is an absolutely necessary step for enterprise deployment. For smaller deployments, this can be a relatively informal process. For larger deployments, it can be quite involved and require a degree of training and experience, and a reasonable probability of the use of highly experienced professionals. For any site survey, software, temporary mounting abilities, measurement tools, and documentation are required. For larger or frequent site surveys, more advanced software and additional tools like attenuators and battery packs are desirable.
- Having deployed an enterprise Wi-Fi infrastructure, there is an ongoing need to maintain and monitor the LAN, and the need for recurring site surveys and ongoing network maintenance, management, and documentation increases proportionally with the scale of the deployment, along with the required level of security. It is in these areas where the differences between lower-cost access points and enterprise access points become particularly apparent. Lower-cost devices typically provide for only browser-based management designed for small-scale deployments, and are not manageable in large deployments because of the requirement to manipulate each device. Enterprise access points provide for management via the CLI, the browser, and NMSs through the SNMP. Enterprise access points are increasingly providing the same features found in wired switches and routers and they better provide for the scalability required in the enterprise.