802.11 Wireless Networks: The Definitive Guide

Creating and Administering Wireless Networks

By Matthew Gast
April 2002
0-596-00183-5, Order Number: 1835
464 pages, $44.95 US $69.95 CA #31.95 UK

802.11 Network Deployment

Deploying a wireless LAN is a considerable undertaking. Significant planning is required before you can even touch the hardware. Deploying a wireless network is not simply a matter of identifying user locations and connecting them to the backbone. Wireless LANs provide mobility through roaming capabilities, but this feature comes with a price. Wireless LANs are much more susceptible to eavesdropping and unauthorized access. Working to mitigate the security problems while offering high levels of service makes large wireless LAN deployments topologically more complex, especially because solving security problems means that a great deal of integration work may be required to get all the different pieces of the solution working in concert.

Wireless networks require far more deployment planning because of the nature of the radio link. Every building has its own personality with respect to radio transmissions, and unexpected interference can pop up nearly everywhere because of microwave ovens, electrical conduits, or severe multipath interference. As a result, each wireless LAN deployment is unique in many respects, and careful planning and a meticulous site survey are required before removing any equipment from the box.

Beyond considerations due to the physical environment, wireless networks often extend an existing wired infrastructure. The wired infrastructure may be quite complex to begin with, especially if it spans several buildings in a campus setting. Wireless networks depend on having a solid, stable, well-designed wired network in place. If the existing network is not stable, chances are the wireless extension is doomed to instability as well.

This chapter is about deployment considerations for wireless LANs, written from a technical perspective. How do the features of wireless LANs influence network topology? Besides the 802.11 equipment, what do you need to deploy a network? How should the logical network be constructed for maximum mobility? What do you need to look for in a site survey to make a deployment successful?

The Topology Archetype

Figure 15-1 shows how many wireless LAN deployments evolve. This figure serves as the road map for this chapter. The guiding principle of Figure 15-1 is that mobility must be limited to the link layer, because network-layer mobility is not generally available on IP networks. The other design decisions help augment the access control of the wireless device and lower management overhead by taking advantage of existing services, each of which will be considered in turn.

Figure 15-1. Standard wireless LAN deployment topology

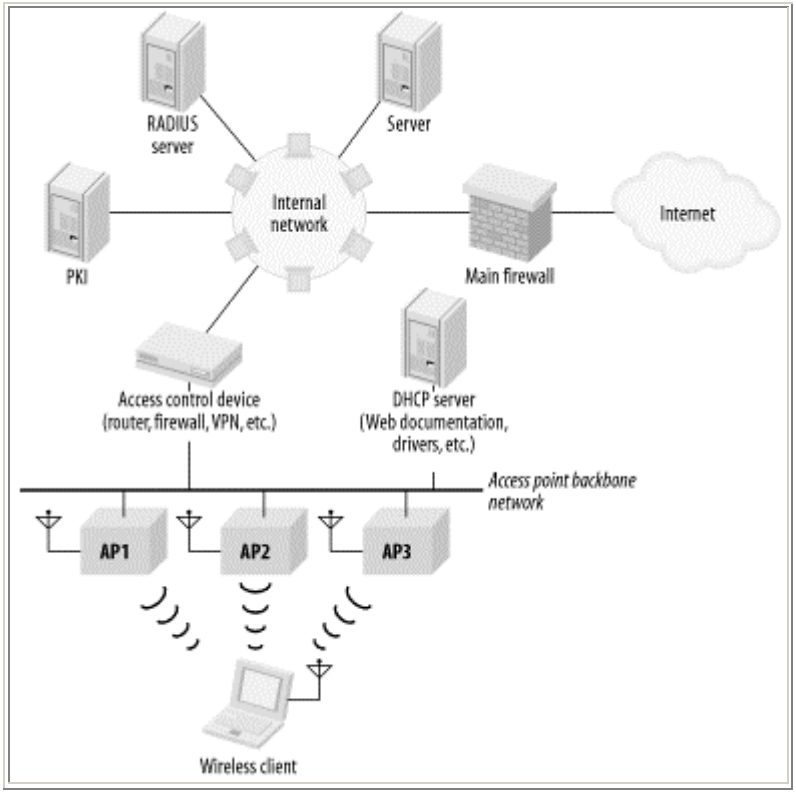

Some deployments may look like multiple instances of Figure 15-1. The topology shown in the figure provides seamless mobility between the access points connected to the access point backbone network. In very large deployments, such as a campus-wide deployment across a large number of buildings, it may be desirable to limit the coverage areas in which seamless roaming is provided. One common strategy is to provide seamless mobility within individual buildings, but not provide roaming between buildings. Each building would have a wireless LAN that looked something like Figure 15-1, and all the access point backbone networks would ultimately connect to a campus backbone.

Roaming and Mobility

In Figure 15-1, the network linking all the access points, which I call the *access point backbone*, is a single IP subnet. To allow users to roam between access points, the network should be a single IP subnet, even if it spans multiple locations, because IP does not generally allow for network-layer mobility. To understand this design

restriction, it is important first to appreciate the difference between true *mobility* and mere *portability*.[1]

Portability certainly results in a net productivity gain because users can access information resources wherever it is convenient to do so. At the core, however, portability removes only the physical barriers to connectivity. It is easy to carry a laptop between several locations, so people do. But portability does not change the ritual of connecting to networks at each new location. It is still necessary to physically connect to the network and reestablish network connections, and network connections cannot be used while the device is being moved.

Mobility, on the other hand, is a far more powerful concept: it removes further barriers, most of which are based on the logical network architecture. Network connections stay active even while the device is in motion. This is critical for tasks requiring persistent, long-lived connections, which may be found in database applications. Support personnel frequently access a tracking database that logs questions, problems, and resolutions. The same argument can be made for a number of tracking applications in a health care setting. Accessing the database through a wireless network can boost productivity because it allows people to add small amounts of information from different locations without needing to reconnect to the database each time. Inventory applications are another example and one of the reasons why retail and logistics are two of the markets that have been quicker to adopt 802.11. When taking inventory, it makes far more sense to count boxes or products where they sit and relay data over a wireless network than to record data on paper and collate the data at the end of the process.

Traditional wired Ethernet connections provide portability. I can take my laptop computer anywhere on the campus at work and plug in. (If I'm willing to tolerate slow speeds, I can even make a phone call and access my corporate network from anywhere in the world.) Each time I access the network, though, I'm starting from scratch. I have to reestablish connections, even if I only moved a few feet. What I'd really like is to walk into the conference room and connect to the corporate network without doing anything.

And therein lies the rub. 802.11 is implemented at the link layer and provides link-layer mobility. IP affords the network designer no such luxury. 802.11 hosts can move within the last network freely, but IP, as it is currently deployed, provides no way to move across subnet boundaries. To the IP-based hosts of the outside world, the VPN/access control boxes of Figure 15-1 are the last-hop routers. To get to an 802.11 wireless station with an IP address on the wireless network, simply go through the IP router to that network. It doesn't matter whether a wireless station is connected to the first or third access point because it is reachable through the last-hop router. As far as the outside world can tell, the wireless station might as well be a workstation connected to an Ethernet.

A second requirement for mobility is that the IP address does not change when connecting to any of the access points. New IP addresses interrupt open connections. If a wireless station connects to the first access point, it must keep the same address when it connects to the third access point.

A corollary to the second requirement is that all the wireless stations must be on the same IP subnet. As long as a station stays on the same IP subnet, it does not need to reinitialize its networking stack and can keep its TCP connections open. If it leaves the subnet, though, it needs to get a IP new address and reestablish any open connections. The purpose of the design in Figure 15-1 is to assign a single IP subnet to the wireless stations and allow them to move freely between access points. Multiple subnets are not forbidden, but if you have different IP subnets, seamless mobility between subnets is not possible.

The "single IP subnet backbone" restriction of the design in Figure 15-1 is a reflection on the technology deployed within most organizations. Mobile IP was standardized in late 1996 in RFC 2002, but it has yet to see widespread deployment. (See the sidebar for a description of how Mobile IP allows stations to change IP addresses without interrupting connections.) Until Mobile IP can be deployed, network designers must live within the limitations of IP and design networks based on fixed locations for IP addresses. In Figure 15-1, the backbone network may be

physically large, but it is fundamentally constrained by the requirement that all access points connect directly to the backbone router (and each other) at the link layer.

Spanning multiple locations with an 802.11 network

Access points that cooperate in providing mobility need to be connected to each other at layer 2. One method of doing this, shown in Figure 15-2a, builds the wireless infrastructure of Figure 15-1 in parallel to the existing wired infrastructure. Access points are supported by a separate set of switches, cables, and uplinks in the core network. Virtual LANs (VLANs) can be employed to cut down on the required physical infrastructure, as in Figure 15-2b. Rather than acting as a simple layer-2 repeater, the switch in Figure 15-2b can logically divide its ports into multiple layer-2 networks. The access points can be placed on a separate VLAN from the existing wired stations, and the "wireless VLAN" can be given its own IP subnet. Frames leaving the switch for the network core are tagged with the VLAN number to keep them logically distinct and may be sent to different destinations based on the tag. Multiple subnets can be run over the same uplink because the VLAN tag allows frames to be logically separated. Incoming frames for the wired networks are tagged with one VLAN identifier, and frames for the wireless VLAN are tagged with a different VLAN identifier. Frames are sent only to ports on the switch that are part of the same VLAN, so incoming frames tagged with the wireless VLAN are delivered only to the access points.

Figure 15-2. Physical topologies for 802.11 network deployment

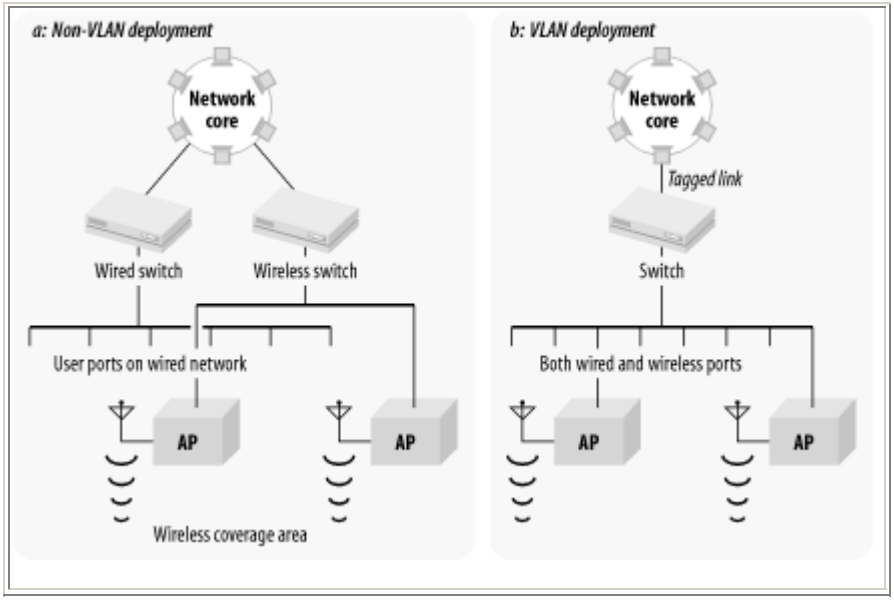

Even better, VLANs can easily span long distances. VLAN-aware switches can be connected to each other, and the tagged link can be used to join multiple physical locations into a single logical network. In Figure 15-3, two switches are connected by a tagged link, and all four access points are assigned to the same VLAN. The four access points can be put on the same IP subnet and will act as if they are connected to a single hub. The tagged link allows the two switches to be separated, and the distance can depend on the technology. By using fiber-optic links, VLANs can be made to go between buildings, so a single IP subnet can be extended across as many buildings as necessary.

Figure 15-3. Using VLANs to span multiple switches

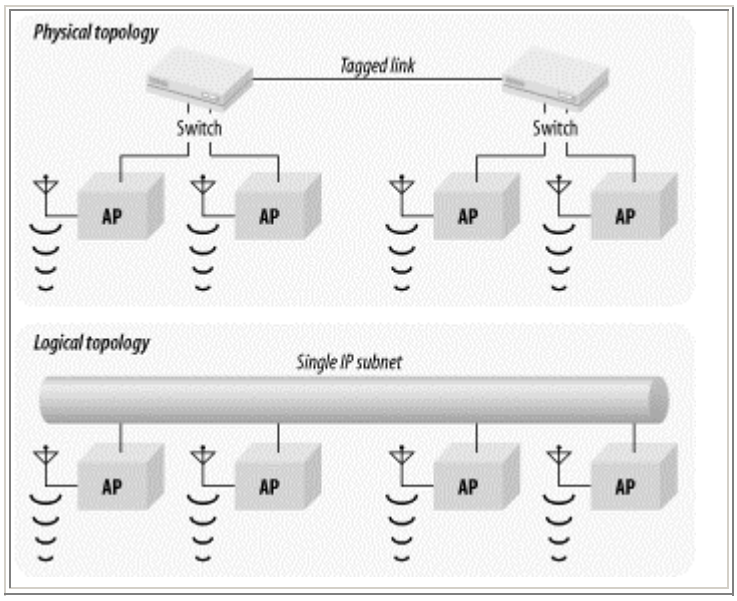

Tagged links can vary widely in cost and complexity. To connect different physical locations in one building, you can use a regular copper Ethernet cable. To connect two buildings together, fiber-optic cable is a must. Different buildings are usually at different voltage levels relative to each other. Connecting two buildings with a conductor such as copper would enable current to flow between (and possibly through) the two Ethernet switches, resulting in expensive damage. Fiber-optic cable does not conduct electricity and will not pick up electrical noise in the outdoor environment, which is a particular concern during electrical storms. Fiber also has the added benefit of high speeds for long-distance transmissions. If several Fast Ethernet devices are connected to a switch, the uplink will be a bottleneck if it is only a Fast Ethernet interface. For best results on larger networks, uplinks are typically Gigabit Ethernet.

For very large organizations with very large budgets, uplinks do not need to be Ethernet. One company I have worked with uses a metro-area ATM cloud to connect buildings throughout a city at the link layer. With appropriate translations between Ethernet and ATM, such a service can be used as a trunk between switches. Computer trade shows such as Comdex and Interop regularly use metro-area networks to showcase both the metro-area services and the equipment used to access those services.

Limits on mobility

The access point backbone network must be a single IP subnet and a single layer-2 connection throughout an area over which continuous coverage is provided. It may span multiple locations using VLANs. Large campuses may be forced to break up the access point backbone network into several smaller networks, each of which resembles Figure 15-1.

802.11 allows an ESS to extend across subnet boundaries, as in Figure 15-4a. Users can roam throughout each "island" of connectivity, but network connections will be interrupted when moving between islands. One solution is to teach users one SSID and let them know that mobility is restricted; another alternative is to name each SSID separately. Both solutions have advantages. In the first case, there is only one SSID and no user confusion, but there may be complaints if the coverage areas do not provide mobility in the right ways. In the second case, mobility is always provided within an SSID, but there are several SSIDs and more opportunity for user confusion.

Figure 15-4. Noncontiguous deployments

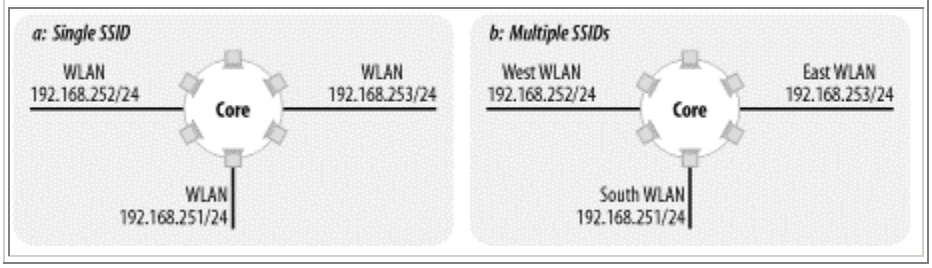

When a campus is broken into several disjointed coverage areas as in Figure 15-4, be sure to preserve the mobility most important to the users. In most cases, mobility within a building will be the most important, so each building's wireless network can be its own IP subnet. In some environments, mobility may be restricted to groups of several buildings each, so the islands in Figure 15-4 may consist of multiple buildings.

Address assignment through DHCP

Multiple independent data sets that must be synchronized are an accident waiting to happen in any field. With respect to wireless LANs, they present a particular problem for DHCP service. To make life as easy as possible for users, it would be best if stations automatically configured themselves with IP network information. DHCP is the best way to do this. Access points frequently include DHCP servers, but it would be folly to activate the DHCP server on every access point. Multiple independent DHCP lease databases are the network equivalent of a tanker-truck pile-up waiting to happen. To avoid the "multiple independent database" problem, have a single source for IP addressing information. Furthermore, some access points may reclaim addresses if an association lapses, regardless of whether the lease has expired. For these reasons, I recommend using an existing DHCP server or installing a new server specifically to support wireless clients. Depending on the importance of the wireless infrastructure, it may be worth considering a backup server as well.

Mobile IP and Roaming

802.11 performs a sleight-of-hand trick with MAC addresses: stations communicate with a MAC address as if it were fixed in place, just like any other Ethernet station. Instead of being fixed in a set location, however, access points note when the mobile station is nearby and relay frames from the wired network to it over the airwaves. It does not matter which access point the mobile station associates with because the appropriate access point performs the relay function. The station on the wired network can communicate with the mobile station as if it were directly attached to the wire.

Mobile IP performs a similar trick with IP addresses. The outside world uses a single IP address that appears to remain in a fixed location, called the *home location*. Rather than being serviced by a user's system, however, the IP address at the home location (the *home address*) is serviced by what is called the *home agent*. Like the access point, the home agent is responsible for keeping track of the current location of the mobile node. When the mobile node is "at home," packets can simply be delivered directly to it. If the mobile node attaches to a different network (called a *foreign network* or *visited network*), it *registers* its so-called foreign location with the home agent so that the home agent can redirect all traffic from the home address to the mobile node on the foreign network.

Consider an example in which two wireless LANs are built on different IP subnets. On its home

subnet, a wireless station can send and receive traffic "normally," since it is on its home network.

When the wireless station moves from its home subnet to the second subnet, it attaches to the network using the normal procedure. It associates with an access point and probably requests an IP address using DHCP. On a wireless station that is unable to use Mobile IP, connections are interrupted at this point because the IP address changes suddenly, invalidating the state of all open TCP connections.

Wireless stations equipped with Mobile IP software, however, can preserve connection state by registering with the home agent. The home agent can accept packets for the mobile station, check its registration tables, and then send the packets to the mobile station at its current location. The mobile station has, in effect, two addresses. It has its home address, and it can continue to use this address for connections that were established using the home address. It may also use the address it has been assigned on the foreign network. No TCP state is invalidated because the mobile station never stopped using its home address.

Naturally, this sidebar has omitted a great deal of the detail of the protocol operations. Designing a protocol to allow a station to attach anywhere in the world and use an address from its home network is a significant engineering endeavor. Several security problems are evident, most notably the authentication of protocol operations and the security of the redirected packets from the home network to the mobile station's current location. Maintaining accurate routing information, both the traditional forwarding tables at Internet gateways and the Mobile IP agents, is a major challenge. And, of course, the protocol must work with both IPv4 and IPv6. For a far more detailed treatment of Mobile IP, I highly recommend *Mobile IP: Design Principles and Practices* by Charles Perkins (Prentice Hall).

Within the context of Figure 15-1, there are two places to put a DHCP server. One is on the access point backbone subnet itself. A standalone DHCP server would be responsible for the addresses available for wireless stations on the wireless subnet. Each subnet would require a DHCP server as part of the rollout. Alternatively, most devices capable of routing also include DHCP relay. The security device shown in Figure 15-1 includes routing capabilities, and many firewalls and VPN devices include DHCP relay. With DHCP relay, requests from the wireless network are

bridged to the access point backbone by the access point and then further relayed by the access controller to the main corporate DHCP server. If your organization centralizes address assignment with DHCP, take advantage of the established, reliable DHCP service by using DHCP relay. One drawback to DHCP relay is that the relay process requires additional time and not all clients will wait patiently, so DHCP relay may not be an option.

Static addressing is acceptable, of course. The drawback to static addressing is that more addresses are required because all users, active or not, are using an address. To minimize end user configuration, it is worth considering using DHCP to assign fixed addresses to MAC addresses.

As a final point, there may be an interaction between address assignment and security. If VPN solutions are deployed, it is possible to use RFC 1918 (private) address space for the infrastructure. DHCP servers could hand out private addresses that enable nodes to reach the VPN servers, and the VPN servers hand out routable addresses once VPN authentication succeeds.

> **TIP:** Use a single DHCP server per access point backbone or DHCP relay at the access point network router to assign addresses to wireless stations. Static addressing or fixed addressing through DHCP is also acceptable.

Security

Informally, data security is defined in terms of three attributes, all of which must be maintained to ensure security:[2]

Integrity
Broadly speaking, integrity is compromised when data is modified by unauthorized users. ("Has somebody improperly changed the data?")
Secrecy
Of the three items, secrecy is perhaps the easiest to understand. We all have secrets and can easily understand the effect of a leak. ("Has the data been improperly disclosed?")
Availability
Data is only as good as your ability to use it. Denial-of-service attacks are the most common threat to availability. ("Can I read my data when I want to?")

Wireless LAN technology has taken a fair number of knocks for its failures in all three areas. Most notably, though, wireless LANs have two major failings with respect to the informal definition of security. First, secrecy is difficult on a wireless network. Wireless networks do not have firm physical boundaries, and frames are transmitted throughout a general area. Attackers can passively listen for frames and analyze data. To defeat attacks against secrecy, network security engineers must employ cryptographic protocols to ensure the confidentiality of data as it travels across the wireless medium. WEP has been a failure in this respect, but other protocols and standards may be employed instead of or in addition to WEP.

Second, integrity may be compromised by wireless hosts. Quick wireless LAN deployments are often connected directly to a supposedly secure internal network, allowing attackers to bypass the firewall. In many institutions, internal stations are afforded higher levels of access privileges. Physical security may have made some additional privileges rational or inevitable, but wireless stations may not necessarily be trusted hosts run by internal users. Attacks against integrity may frequently be defeated by strong access control.

Vendors often tout WEP as a security solution, but the proven flaws in the design of WEP should give even the most freewheeling security administrators cause for concern. WEP is, in the words of one industry observer, "unsafe at any key length."[3] Future approaches based on 802.1x and EAP may improve the picture, but current deployments must depend on solutions that are available now. Although products claiming to support 802.1x are currently appearing on the market, they have yet to establish a track record with respect to either security or interoperabilty.

Access control and authentication

Connecting to wireless networks is designed to be easy. In fact, the ease of connection is one of the major advantages to many newer wireless technologies. 802.11 networks announce themselves to anybody willing to listen for the Beacon frames, and access control is limited by the primitive tools supplied by 802.11 itself. To protect networks against the threat of unauthorized access, strong access control

should be applied. A helpful rule of thumb is to treat wireless access points like open network drops in the building lobby. 802.11 networks can benefit from access control at two points:

- Before associating with an access point, wireless stations must first authenticate. At present, this process is either nonexistent or based on WEP.
- After association with the access point, the wireless station is attached to the wireless network. However, strong authentication can be applied to any wireless stations to ensure that only authorized users are connecting to protected resources. This form of access control is no different from the access control widely enforced by firewalls today.

At the present time, the initial authentication during the association process is pitifully weak. Current deployments must depend on two methods, one of which was never specified by the standard but is widely used.

One approach is to allow only a specified set of wireless LAN interface MAC addresses to connect to access points. Maintaining the list is its own administrative headache. Distributing the list to access points may be even worse. In a network with access points from multiple vendors, the script may need to massage the list into different file formats to cope with what different products require. Frequently, the list of allowed devices must be distributed by TFTP. Even if the distribution is automated by administrative scripts, TFTP comes with its own security woes. Furthermore, like wired Ethernet cards, 802.11 cards may change the transmitter MAC address, which totally undermines the use of the MAC address as an access control token. Attackers equipped with packet sniffers can easily monitor successful associations to acquire a list of allowed MAC addresses.

A second approach is to allow connections from stations that possess a valid WEP key. Stations that pass the WEP challenge are associated, and stations that fail are not. As noted in Chapter 5, this method is not very strong because WEP is based on RC4, and it is possible to fake a legitimate response to a WEP challenge without any knowledge of the WEP key. In spite of its limitations, WEP makes a useful speed

bump for attackers to jump over. Use it, but be aware of its limitations. Or disable it, but be cognizant of the fact that association is unrestricted.

In some products, these methods may be combined. However, both are easily defeated. Maintaining strong security over a wireless LAN requires solutions outside the scope of 802.11, in large part to augment the relatively weak access control supplied by 802.11.

Many networks deploy firewalls to protect against unauthorized access and use of systems by outsiders. In many respects, wireless stations should be considered untrusted until they prove otherwise, simply because of the lack of control over the physical connection. In the network topology shown in Figure 15-1, an access control device is used to protect the internal network from wireless stations. This access control device could be one of several things: a firewall, a VPN termination device, or a custom solution tailored to the requirements of 802.11 networks.

At the time this book was written, many security-conscious organizations opted to use existing firewalls or VPN devices or build systems to meet their own internal requirements. Firewalls are well-known for providing a number of strong authentication mechanisms, and they have a proven ability to integrate with one-time password systems such as RSA's SecurID tokens. New releases of IPSec VPN devices also increasingly have this capability. Initial versions of the IPSec specification allowed authentication only through digital certificates. Certificates work well for site-to-site VPNs, but the idea of rolling out a public-key infrastructure (PKI) to support remote access was frightening for most users. As a result, several new approaches allow for traditional ("legacy") user authentication mechanisms by passing VPN user authentication requests to a RADIUS server. Several mechanisms were in draft form as this book was written: Extended Authentication (XAUTH), Hybrid Mode IKE, and CRACK (Challenge/Response for Authenticated Control Keys).

Several wireless LAN vendors have also stepped up to the plate to offer specialized "wireless access controller" devices, which typically combine packet filtering,

authentication, authorization, and accounting services (AAA), and a DHCP server; many devices also include a DNS server and VPN termination. AAA features are typically provided by an interface to an existing corporate infrastructure such as RADIUS, which frequently has already been configured for remote access purposes. Some products may also include dynamic DNS so that a domain name is assigned to a user, but the IP number can be assigned with DHCP.

Several vendors have access controller solutions. Cisco offers an external access control server for the Aironet product line. Lucent's ORiNOCO AS-2000 access server has an integrated RADIUS server. Nokia's P020 Public Access Zone Controller is an integrated network appliance with a RADIUS client and DHCP server, and the companion P030 Mobility Services Manager offers the RADIUS server and billing functions.

> **TIP:** Recognize the limitations of WEP. Treat wireless stations as you would treat untrusted external hosts. Isolate wireless LAN segments with firewalls, and use strong authentication for access control. Consider using existing user databases as part of the authentication roll-out.

Confidentiality: WEP, IPSec, or something else?

Confidentiality is the second major goal in wireless LAN deployments. Traffic is left unprotected by default, and this is an inappropriate security posture for most organizations. Users can choose among three options:

- Use WEP.
- Use a proven cryptographic product based on open protocols.
- Use a proprietary protocol.

Option three locks you into a single vendor and leaves you at their mercy for upgrades and bug fixes. Proprietary cryptographic protocols also have a poor track record at ensuring security. In the end, the choice really comes down to whether WEP is good enough. Given the insecurity of WEP, there are two questions to ask:

"Does the data on this network need to stay secret for more than a week?"

- WEP is not strong encryption by any stretch of the imagination, and you should assume that a sufficiently motivated attacker could easily capture traffic from the wireless network, recover the WEP key, and decrypt any data.

"Do users need to be protected from each other?"

- In most WEP deployments, keys are distributed to every authorized station. When all users have access to the key, the data is protected from outsiders only. WEP does not protect an authorized user with the key from recovering the data transmitted by another authorized user. If users need to be protected from each other, which is a common requirement in many computing environments, then additional security precautions are required.

Choosing a cryptographic protocol or product is subject to a few basic ground rules, conveniently summarized in the Cryptographic Snake Oil FAQ.[4] While looking for signs of snake oil is not sure protection, it should filter out the most egregious duds. Cryptography is like a fine wine--it gets better with age. If a protocol or algorithm has withstood extensive public analysis, it is probably better than something just invented.

There are only a few non-snake oil solutions that are worth considering. To provide confidentiality at the network layer, there is only one standard: IPSec. Unfortunately, IPSec is not without its drawbacks. It is a complex system to understand and use, so you must be prepared for a learning curve for network administrators. The complexity of IPSec contributes to a relatively high management overhead, at least at the beginning of deployment. IPSec solutions require the installation of client software on wireless stations to protect outbound traffic, and desktop software management is always unpleasant at best. Perhaps the most frustrating attribute of IPSec is the difficulty in configuring two different systems to be interoperable. Extensive testing of IPSec interoperability can be a huge burden, one that should not be taken lightly.

An alternative is to allow only applications with strong built-in cryptographic systems. Web-based systems can be secured with the secure socket layer (SSL). Host logins can be secured with SSH. SSH can also be used to secure many types of TCP-based network traffic, though the port-forwarding configuration may be too complex for many users. Some environments may have already deployed a framework such as Kerberos for application layer security, in which case, it can probably be extended to wireless stations without great difficulty.

> **TIP:** Consider the wireless network to be transmitting data in the clear if you are not using strong, proven cryptographic solutions such as IPSec or SSH.

Availability through redundancy

Nothing in Figure 15-1 requires single points of failure. Clustered solutions for all the major components exist, so availability is not necessarily compromised by the failure of any of the security components. Clusters are composed of several independent devices that get together and share state information among multiple machines. When any member of the cluster fails, survivors pick up the workload with no interruption. If you elect to use either firewalls or VPNs, consider using a clustered product. Clustering is particularly important for VPN access because you own the infrastructure, and users are far less forgiving of internal problems than flaky Internet connectivity.

One item to watch for in this area is redundancy for DHCP servers. No standard exists for synchronizing the data held by DHCP servers. However, the Network Working Group of the IETF is working towards a standardized DHCP failover protocol, which will increase the reliability of the address allocation service.

Summary and Analysis of Archetypal Topology

Several points about the archetypal topology of Figure 15-1 are important and bear repeating:

1. 802.11 provides for mobility only within an extended service set, and Inter-Access Point Protocols (IAPPs) cooperate only when directly connected at the link layer. Until a standard is developed, proprietary IAPPs are not guaranteed to interoperate; you need to select a single access point vendor for each area of continuous coverage. Each extended service set (ESS) should therefore be a single IP subnet. It is acceptable to use VLANs and other bridging technologies to achieve this goal.
2. Address assignment is best done as dynamic addressing to minimize end user configuration. Only one DHCP server should be responsible for handing out addresses to wireless stations because it is important to prevent accidental readdressing. That server may be placed on the access point backbone. DHCP relay can also be used to take advantage of DHCP servers that are already deployed.
3. Consider the use of WEP. In many cases, it is not recommended because it does not significantly bolster security, but it may complicate roaming, and it is just another configuration item to get wrong. Some vendors charge more for cards that implement the "strong" 128-bit WEP, too.
4. Consider the security policy and goals for your wireless network. WEP is problematic, but it may be better than running open access points. In many environments, though, WEP should be disabled. With large numbers of users, WEP is just another configuration item to get wrong. Deploying WEP may also complicate roaming between access points.
5. Carefully consider the limitations of WEP and deploy additional solutions to enhance:

Authentication

WEP does not provide strong user authentication, so treat any wireless stations as you would untrusted external hosts. Your security policy may offer guidance here--how do you protect against threats from existing mobile users, such as telecommuters and road warriors? Isolate wireless LAN segments with firewalls and use strong authentication for access control. Existing user databases such as RADIUS servers can probably be used as part of the access control deployment.

Confidentiality

If you do not want to broadcast data to the world, use strong, proven cryptographic solutions to protect data as it traverses the airwaves. IPSec is the best standardized solution, especially since it now provides strong user authentication with existing token-based systems.

6. To maintain availability and increase uptime, use clustered solutions whenever possible and cost-effective.

Project Planning

Site survey work is the heart of installing a wireless LAN. To successfully run a site survey, though, preparation is very important.

Gathering Requirements

Before "breaking cable" on a network expansion, gather end user requirements and information to find out which expectations are important. Use the following checklist to flesh out the customer requirements; each point is detailed further in a subsequent section:

Throughput considerations

How much throughput is required? This is partly dependent on the type of device that will be used on the wireless LAN, though if it is a PC-like device with the ability to display large and complex graphics, you will want your wireless LAN to be as fast as possible. In most cases, this will lead to choosing 802.11b-based networks to use the 11-Mbps physical layer. If you like leading-edge technology, you may want to consider 802.11a products, several of which appeared just as this book went to press.

Coverage area

Where should coverage be provided, and what is the density of users in particular areas?

Mobility

How much movement between coverage is needed? Does it need to be full mobility, with continuous connections as the wireless station moves around the network? Or can the network simply enable effective portability by facilitating automatic reconfiguration when the mobile station moves between coverage areas?

User population

How many people will use the wireless network, and what quality of service do they expect? As always, be sure to allow for growth!

Physical network planning

Will new network cabling be needed to supply the wireless LAN backbone, or can you make do with existing cabling? Are an adequate number of outlets available in the correct locations? Can the access points and antennas be installed in the open or must they be confined to wiring closets or other hidden locations?

Logical network planning

How many IP addresses will be set aside for wireless users? Is a large enough address block available, or will the existing network need to be renumbered to accommodate the wireless network? If the necessary IP address space is not available, it may mean cutting back on the level of seamless mobility on the wireless LAN.

Application characteristics

Can address translation be used to save IP address space? Are any applications sensitive to high or variable delays? Do any applications provide time-critical data? If so, consider looking for products that support the point coordination function and contention-free delivery, but be aware that many products do not support the PCF.

Security requirements

Wireless LANs have been subject to a number of security concerns. The two main goals of wireless LAN security planning are ensuring adequate access control and preserving the confidentiality of data as it traverses the wireless network. Security requirements may be dictated by legal requirements or the legal threat of unauthorized data disclosure.

Authentication has long been a weak point of 802.11 networks. The two main options provided by 802.11 are to filter on the MAC addresses of wireless stations allowed to connect or use shared WEP keys for stronger authentication. In practice, MAC address filtering is too cumbersome and error-prone, so the choices are to use WEP authentication or depend on external solutions.

Data confidentiality is provided by encryption services. One option is the WEP standard, though higher-security sites may opt for additional VPN technology on top of the 802.11 layer.

Site environmental considerations

A number of factors can affect radio propagation and signal quality. Building materials, construction, and floor plan all affect how well radio waves can move throughout the building. Interference is a fact of life, but it is more pronounced in some buildings than in others. Temperature and humidity have minor effects. Early site visits can assist in anticipating several factors, and a detailed site survey can spot any real problems before installation begins in earnest.

Purchasing wireless LAN hardware and software

At some point, wireless LAN hardware and software must be purchased. Many vendors exist, and the decision can be based on a number of criteria. (Selecting an access point vendor was discussed in Chapter 14.) Selecting cards may depend on your institution's policies. Some organizations may choose to purchase all cards centrally from a single vendor. Others may choose to select a small set of officially "supported" vendors but allow users to select alternative hardware if they are willing to forego official support from the network staff. At least one wireless network analyzer should be part of the budget. Depending on the size of the wireless LAN and the number of network administrators, you may wish to budget for more than one.

Project management

As with many other projects, drawing up a schedule and budget is a necessary component. This chapter does not provide any guidance on nontechnical factors because they are often organization-specific.

Network performance requirements

Depending on the applications used on the wireless network, different requirements are imposed. One of the most important items, and the one that is least under the control of the network architect, is the characteristics of the application. Most applications can now be run over TCP/IP, but they may require widely varying throughput, delay, or timing characteristics. More importantly, though, is how an application reacts to network address translation (NAT)--translating its IP addresses with intermediate devices.

Single TCP connections, such as those used by HTTP and SSH, are easily translated with no side effects. Other network protocols, most notably those in the Microsoft Networking family, embed the source IP address in the data portion of the packet and cannot be used in conjunction with address translation without great difficulty.[5] NAT also causes problems for most videoconferencing applications. At the time this book was written, standardized IPSec also did not work when the IPSec packet was passed through an address translator because IPSec authenticates the source IP address of the packet. Translating the source IP address caused the integrity check to fail.

The remaining three factors are under direct control of the end user. A coverage area must be defined, and a form of mobility between the coverage areas is a likely companion requirement. (Mobility imposes its own requirements on the IP addressing architecture, which was discussed previously.)

Finally, end users will have a target throughput requirement. Any throughput goals must be carefully considered because wireless LANs are a shared medium, but one without an upgrade path similar to dropping in an Ethernet switch.

Table 15-1 shows the number of users that can be served on 11-Mbps 802.11b networks, with different sustained loads per user. Wireless LAN bit rates are low, and the extra management features limit throughput to a relatively low fraction of the available bit rate. As you can see from the table, though, 11-Mbps networks are likely to be practical for office environments, which are mainly email, web browsing, and intermittent file access. It is likely that 20-30 users per access point is a reasonable estimate for capacity planning.

Table 15-1: Network capacity compared to sustained throughput per user

Connection method and speed	Effective number of simultaneous users on 11-Mbps networks (6 Mbps data throughput)
Cellular modem, 9.6 kbps	625
Modem, 50 kbps	120
Single ISDN B channel, 64	93

kbps	
Dual ISDN, 128 kbps	46
100 kbps sustained LAN usage	60
150 kbps sustained LAN usage	40
200 kbps sustained LAN usage	30
300 kbps sustained LAN usage	20

Realistic throughput expectations for 802.11b networks

In terms of throughput, the performance of 802.11 LANs is similar to shared Ethernet. As more users are added, the available capacity per user is divided up. A practical rule of thumb is that the highest throughput that can be attained using the DCF is 75% of the nominal bit rate. The 75% figure is a theoretical result derived from the protocol itself; it includes overhead such as the preamble, interframe spaces, and framing headers. However, throughput rates as low as 50% may be observed. A target of 65% of the nominal bit rate is commonly observed.

For 2-Mbps networks, this translates to a top speed of 1.5 Mbps, though rates as low as 1.3 Mbps are common. Applying similar percentages to 11-Mbps networks yields a practical throughput range of 6 to 8 Mbps.

For networks under the operation of the PCF, throughput is higher because it uses shorter interframe spaces and more efficient acknowledgments. Implementation of the PCF is not required by the standard, so implementations are quite uncommon.

11a, 11b, 11g, and more?

Where do you go past 11 Mbps? That's the question people have been asking with increasing frequency over the past year. Right now, the leading standard for higher wireless data rates is 802.11a; it provides 54 Mbps in the 5-GHz band. 802.11a products are on the market now, though it's really too early to say anything substantial about them. Some vendors are announcing that their access points can be upgraded to 802.11a by purchasing a new card and installing new firmware. Software upgrades may be helpful, but only if the hardware is ready for 802.11a. Some

vendors have boasted about the easy software upgrade to the 54-Mbps performance of 802.11a, but the access point in question had only a 10-Mbps Ethernet port. Any access points you consider software upgrades for should have Fast Ethernet ports. 802.11a is somewhat more expensive than 11b, though prices should begin to drop soon.

Because 802.11a uses the same MAC layer as 802.11b, with the OFDM PHY layer discussed in Chapter 12, I expect the installation and administration of products to be essentially the same as it is for 802.11b products. In short, just about everything discussed in this book still applies. I am guessing that the range should be similar; OFDM looks like a superior modulation technique, but at the higher frequency, there should be greater problems with path loss, multipath fading, reflections, etc. One estimation was that the radius of 802.11a access points would be 20-25% shorter. Because the 5GHz band is much larger than the ISM band and isn't already occupied by microwave ovens and other devices, there should be fewer problems with interference.

Another standard waiting in the wings is 802.11g. 11g is a 2.4-GHz standard, like 11b, but it uses the OFDM modulation technique of 11a. It also operates at 54 Mbps. The standard isn't finalized yet, and it's hard to imagine products appearing before the end of 2002. The upgrade path from 11b to 11g might be easier than that from 11b to 11a; because both standards use the same frequency band, you should be able to upgrade your access points without worrying about changing their coverage. (For better or worse, the RF characteristics of your site will be different at 2.4 and 5 GHz, so you may find you need to move or add access points if you migrate to the higher-frequency band.) 802.11g also promises to be less expensive than 11a, though in practice this probably means only that the existence of 11g will drive down the price of 11a products.

Security

The security trade-offs were discussed in the previous "Security" section. In many cases, an IPSec-based VPN is the logical choice. IPSec was designed for precisely the

environment that wireless LANs typify. Intruders can easily capture traffic and perform extensive offline attacks with stored data. Now that IPSec is evolving to support remote clients connecting to central sites, it can also be used to provide strong authentication without a difficult PKI rollout. Many products now support one of the various standards to allow an IPSec termination device to perform user authentication through RADIUS, which allows administrators to take advantage of existing authentication databases. The new 802.1x standard incorporates RADIUS; unless there's a critical problem in 802.1x (as there was in WEP), future wireless security is likely to be based on the 1x standard.

Coverage and physical installation restrictions

Part of the end user requirement is a desired coverage area, and possibly some physical restrictions to go along with it. Physical restrictions, such as a lack of available electrical power and network connections, can be mundane. Some institutions may also require that access points and antennas are hidden; this may be to maintain the physical security of the network infrastructure, or it may be simply to preserve the aesthetic appeal of the building.

Some organizations may want to provide coverage outdoors as well, though this is confined to mild climates. Any equipment placed outdoors should be sturdy enough to work there, which is largely a matter of waterproofing and weather resistance. One solution is to install access points inside and run antennas to outdoor locations, but external antenna cables that are long enough are not always available. Outdoor network extensions can be difficult because most 802.11 equipment is not suited to outdoor use, and even if it was, power and Ethernet connections are not readily available outdoors. The best approach to providing outdoor coverage is to keep the access points inside and use external weatherproof antennas on the roof.

The Building

It is a great help to get blueprints or floor plans and take a tour of the installation site as early as possible in the process. Based on a walk-through with the floor plans, you can note where coverage must be provided, nearby network and power drops, and

any relevant environmental factors. Most importantly, you can correct the blueprints based on any changes made to the structure since the blueprints were drawn. Many minor changes will not be reflected on the blueprints.

Different materials have different effects on the radio link. Signal power is most affected by metal, so elevator shafts and air ducts cause significant disruption of communications. Tinted or coated windows frequently cause severe disruption of radio signals. Some buildings may have metal-coated ceilings or significant amounts of metal in the floor. Wood and most glass panes have only small effects, though bulletproof glass can be quite bad. Brick and concrete have effects somewhere between metal and plain untreated glass. To a large extent, though, the expected drop in signal quality due to building construction is a judgment call that improves with experience.

During a pre-survey walk-though, also note any potential sources of interference. The 2.4-GHz ISM band is unlicensed, so many types of devices using the band can be deployed without central coordination. Newer cordless phones operate in the 2.4-GHz band, as well as Bluetooth-based devices and a number of other unlicensed radio devices. Depending on the quality and amount of shielding, microwave ovens may also emit enough radiation to disrupt 802.11 communications. If you anticipate a large amount of interference, testing tools called *spectrum analyzers* can identify the amount of radiation in the wireless LAN frequency band. If your organization does RF testing, it may be necessary to shield any labs where testing is done to avoid interference with the wireless LAN. As a rule of thumb, keep access points at least 25 feet away from any strong interference sources. End user devices also suffer if they are located too close to sources of interference, but only end user communications are interrupted in that case.

The Network

There are two components to network planning. The first, physical planning, is largely legwork. In addition to the building map, it helps to obtain a physical network map, if one exists. It is much easier to install wireless LAN hardware when no

expensive and time-consuming wiring needs to be done. Knowing the location and contents of all the wiring closets is an important first step.

The second component of network planning is the plan for changes to the logical network. How will mobile stations be addressed? How will access points be reconnected to their firewall or router?

Network addressing

802.11 provides mobility between access points as long as both access points are part of the same ESS. Roaming works only when mobile stations can transfer from one access point to another and keep the same IP address. All access points of the same ESS must therefore be connected to the same IP subnet so that wireless stations can keep addresses as they associate with different access points.

To get the IP space allocated, you will probably need to work with a network administrator. The administrator will want to know how many addresses you need and why. In addition to the planned number of wireless stations, be sure to include an address for each access point and any servers and security devices on the wireless subnet.

After tentatively locating access points on a blueprint or sketch of the area, work with the physical network map to plug the access points into the nearest wiring closet. If the access device is a switch with VLAN capability, the access point can probably be placed on the access point backbone VLAN. If not, it may be necessary to patch the access point back to a switch capable of VLAN connections or replace the access device with a small multi-VLAN switch.

Preliminary Plan

Based on the floor plans, use the map to come up with a preliminary plan. The preliminary plan is based on the coverage area required and the typical coverage radius from an access point. At this point, detailed radio channel use planning is not yet necessary. The main use of the preliminary plan is to come up with trial access

point locations to begin signal quality measurements as part of the detailed site survey. Table 15-2 is based on the best-case coverage radius from a typical omnidirectional antenna.

Table 15-2: "Rule-of-thumb" coverage radius for different types of space

Type of space	Maximum coverage radius
Closed office	up to 50-60 feet
Open office (cubicles)	up to 90 feet
Hallways and other large rooms	up to 150 feet
Outdoors	up to 300 feet

The Site Survey

After coming up with the preliminary plan, it is time to move on to the heart of the deployment routine: the site survey. Several options exist for performing a site survey. Vendors may provide site surveys to early adopters who agree to be reference accounts. Value-added resellers may also have the skills to perform detailed site surveys; resellers may sell site survey consulting services or use site surveys as a way of coming up with a wireless LAN deployment bid. Some companies that specialize in technical education also offer classes on performing site surveys.

Refining the preliminary design is the purpose of the site survey. Radio transmission is complicated, and some things must be done by experiment. All sites will require adjustments to the preliminary design as part of the site survey. In many cases, use of site survey tools can help eliminate access points from a network design and result in substantial cost savings. The major goal of a site survey is to discover any unforeseen interference and redesign the network accordingly. In most cases, interference problems can be repaired by relocating an access point or using a different antenna.

The site survey should assess the following:

- The actual coverage of the access points and the optimum location of access points in the final network
- Actual bit rates and error rates in different locations, especially locations with a large number of users

- Whether the number of access points is sufficient--more or fewer may be required, depending on the characteristics of the building with respect to radio waves
- The performance characteristics of customer applications on the wireless LAN

Tools

Site survey work consists mostly of seemingly endless signal quality measurements. Depending on the tool used, the signal quality measurements may be any of the following:

Packet Error Rate (PER)

The fraction of frames received in error, without regard to retransmissions. A common rule of thumb is that the PER should be less than 8% for acceptable performance.

Received Signal Strength Indication (RSSI)

A value derived from the underlying mathematics. Higher values correspond to stronger (and presumably better) signals.

Multipath time dispersion

Some software or instruments may be able to measure the degree to which a signal is spread out in time by path differences. Higher delay spreads make the correlation of the wideband signals more difficult. Devices need to accept either a higher error rate at high delay spreads or fall back to a more conservative coding method. Either way, throughput goes down. The higher the delay spread, the more throughput suffers.

Signal quality measurements can be carried out by a dedicated hardware device or a software program running on a laptop with the card vendor's site survey tool. Several wireless LAN vendors, such as Intel, Proxim, and 3Com, bundle site survey tools with their access points. Handheld site survey tools designed specifically for 802.11 networks also exist.

Patience and comfortable shoes are among the most important items to bring to a site survey. Measuring signal quality in an area is a painstaking process, requiring many measurements, often taken after making minor changes to the antenna or

access point configuration. You will spend a lot of time walking, so wear shoes that you can walk all day in.

Particularly stubborn interference may require the use of a *spectrum analyzer* to locate the source of interference from a non-802.11 network. Devices that can scan a wide frequency band to locate transmissions are not cheap. Expect to pay several thousand dollars, or you can hire a consultant. In either case, a spectrum analyzer is the tool of last resort, necessary for only the most stubborn problems.

Antenna Types

Wireless cards all have built-in antennas, but these antennas are, at best, minimally adequate. If you were planning to cover an office--or an even larger area, such as a campus--you will almost certainly want to use external antennas for your access points. When considering specialized antennas, there are only a few specifications that you need to pay attention to:

Antenna type

The antenna type determines its radiation pattern--is it omnidirectional, bidirectional, or unidirectional? Omnidirectional antennas are good for covering large areas; bidirectional antennas are particularly good at covering corridors; unidirectional antennas are best at setting up point-to-point links between buildings, or even different sites.

Gain

The gain of the antenna is the extent to which it enhances the signal in its preferred direction. Antenna gain is measured in dBi, which stands for decibels relative to an isotropic radiator. An isotropic radiator is a theoretical beast that radiates equally in all directions. To put some stakes in the ground: I've never seen a specification for the gain of the built-in antenna on a wireless card, but I would guess that it's negative (i.e., worse than an isotropic radiator). Simple external antennas typically have gains of 3 to 7 dBi. Directional antennas can have gains as high as 24 dBi.[6]

Half-power beam width

This is the width of the antenna's radiation pattern, measured in terms of the points at which the antenna's radiation drops to half of its peak value. Understanding the half-power beam width is important to understanding your antenna's effective coverage area. For a very high-gain antenna, the half-power beam width may be only a couple of degrees. Once you get outside the half-power beam width, the signal typically drops off fairly quickly, though that depends on the antenna's design. Don't be fooled into thinking that the half-power beam width is irrelevant for an omnidirectional antenna. A typical omnidirectional (vertical) antenna is only omnidirectional in the horizontal plane. As you go above or below the plane on which the antenna is mounted, the signal decreases.

We've discussed antennas entirely in terms of their properties for transmitting, largely because most people find that easier to understand. Fortunately, an antenna's receiving properties are identical to its transmitting properties--an antenna enhances a received signal to the same extent that it enhances the transmitted signal. This result is probably what you would expect, but proving it is beyond the scope of this book.

Now, let's talk about some of the antenna types that are available; Figure 15-5 shows a number of different antenna types:

Figure 15-5. Antenna types

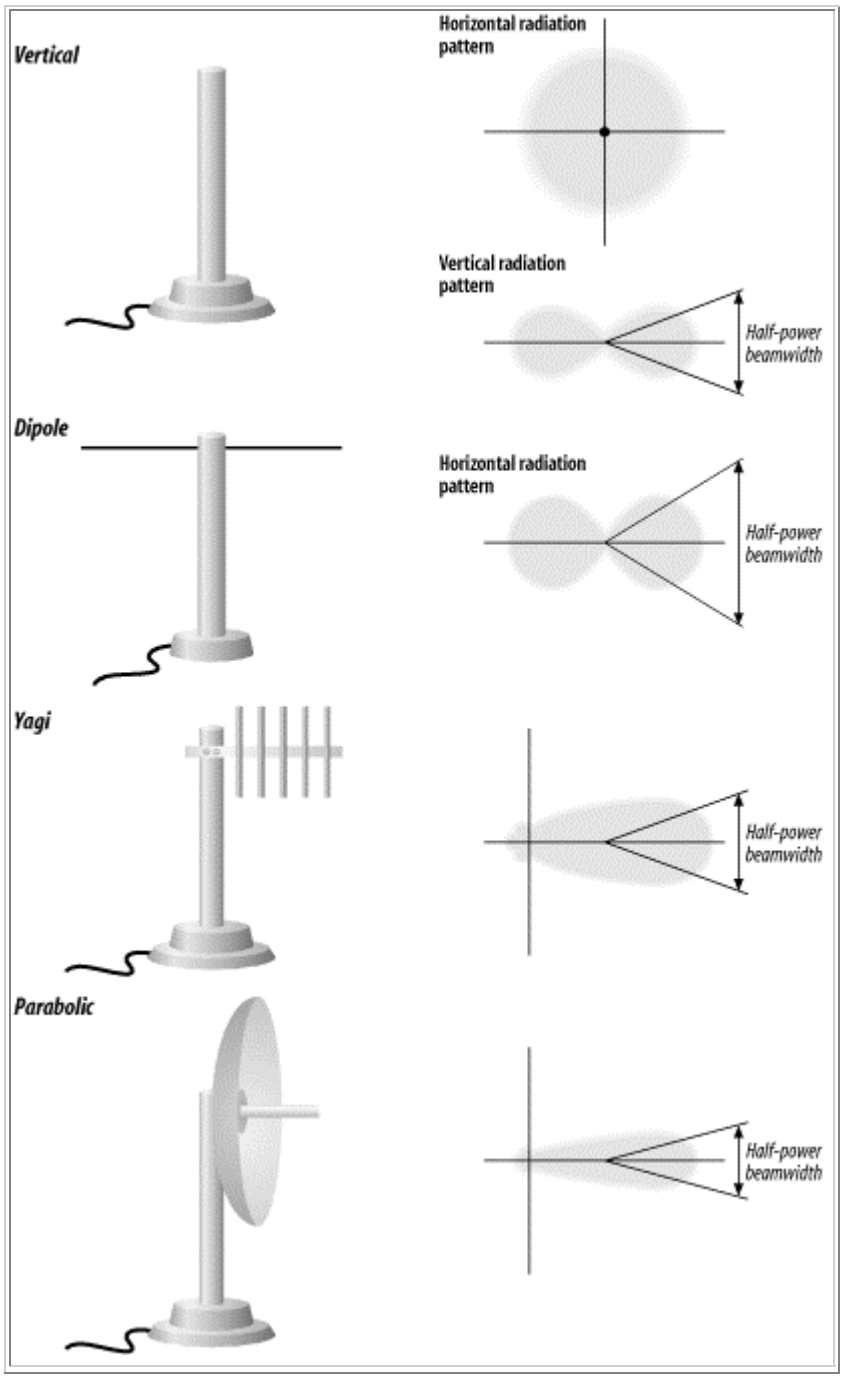

Vertical

This is a garden variety omnidirectional antenna. Most vendors sell several different types of vertical antenna, differing primarily in their gain; you might see a vertical antenna with a published gain as high as 10 dBi or as low as 3 dBi. How does an omnidirectional antenna generate gain? Remember that a vertical antenna is omnidirectional only in the horizontal plane. In three dimensions, its radiation pattern looks something like a donut. A higher gain means that the donut is

squashed. It also means that the antenna is larger and more expensive, though no antennas for 802.11 service are particularly large.

If you want to cover a confined outdoor area--for example, a courtyard between several buildings of a corporate campus--note that the half-power beam width means that a roof-mounted vertical antenna might be less than ideal, particularly if the building is tall. Vertical antennas are good at radiating out horizontally; they're not good at radiating down. In a situation like this, you would be better off mounting the antenna outside a first- or second-story window.

Dipole

A dipole antenna has a figure eight radiation pattern, which means it's ideal for covering a hallway or some other long, thin area. Physically, it won't look much different from a vertical--in fact, some vertical antennas are simply vertically mounted dipoles.

Yagi

A Yagi antenna is a moderately high-gain unidirectional antenna. It looks somewhat like a classic TV antenna. There are a number of parallel metal elements at right angles to a boom. However, you are not likely to see the elements on a Yagi for 802.11 service; the commercially made Yagis that I have seen are all enclosed in a *radome*, which is a plastic shell that protects the antenna from the elements in outdoor deployments. Yagi antennas for 802.11 service have gains between 12 and 18 dBi; aiming them is not as difficult as aiming a parabolic antenna, though it can be tricky.

Parabolic

This is a very high-gain antenna. Because parabolic antennas have very high gains (up to 24 dBi for commercially made 802.11 antennas), they also have very narrow beam widths. You would probably use a parabolic antenna only for a link between buildings; because of the narrow beam width, they are not very useful for providing services to end users. Vendors publish ranges of up to 20 miles for their parabolic antennas. Presumably, both ends of the link are using a similar antenna. Do not underestimate the difficulty of aiming a parabolic antenna properly--one commercial product has a published beam width of only 6.5 degrees. If you decide to install a parabolic antenna, make sure that you have it mounted firmly. You do not want a bad storm to nudge it a bit and take down your connection.

Some vendors make an issue of the distinction between "mesh" or "grid" parabolas (in which the antenna's reflector looks like a bent barbecue grill) and solid parabolas. Don't sweat it--if the antenna is well-designed, the difference in performance between a mesh and a solid reflector is not worth worrying about. A mesh does have an advantage, though, in areas subject to high winds.

Parabolic and Yagi antennas are useful primarily for links between buildings. The biggest problem is aiming them properly. If the two sites are visible to each other, you can play some tricks with gunsights--though if you can see one site from the other, you probably don't need such a sophisticated antenna system. Otherwise, buy a good compass and a topographical map from the U.S. Geological Survey, and compute the heading from one site to the other. Remember to correct for magnetic north. If you can spend some extra money, you might be able to simplify the setup by installing a high-gain vertical antenna at one site; then you need to aim only one antenna. If the signal is marginal, replace the vertical with a parabolic antenna once you have the first antenna aimed correctly.

High-gain antennas can become a regulatory problem, particularly in Europe (where power limits are lower than in the U.S.). Lucent notes that their high-gain parabolic antenna cannot be used legally on channels 1, 2, 10, and 11 in the U.S., though it can be used on channels 3 through 9. But that limitation probably assumes that you're using a Lucent wireless card, and Lucent's transceiver produces less output power than the Intersil chip set used by most other vendors. If you connect the Lucent parabolic antenna to a Nokia wireless card, you'll be be way beyond the maximum legal effective radiated power.

Cabling

Having put so much effort into thinking about antennas, we have to spend some time thinking about how to connect the antennas to the access points or wireless cards. Most vendors sell two kinds of cable: relatively inexpensive thin cable (typically 0.1 inch in diameter) and "low-loss cable" that's substantially thicker (typically 0.4 inch) and much more expensive. The thin cable is usually available only

in lengths of a couple of feet, and that's as it should be: it is very lossy, and more than a few feet can easily eat up your entire signal. It's intended for connecting a wireless card in a laptop to a portable antenna on your desktop, and that's all. To put numbers behind this: one vendor specifies a loss of 2.5 dB for a 2-meter cable. That means that close to half of your signal strength is disappearing in just two meters of cable. One cable vendor, for a cable that would typically be used in this application, specifies a loss of 75 dB per 100 feet at 2.4 GHz. That means that your signal strength will drop by a factor of *225* (roughly 33 million), clearly not something you want to contemplate. I know of one vendor that recommends using RG58 cable with medium-gain antennas. RG58 is better than the really thin cable intended for portable use, but not much better (35 dB per 100 feet); if you use RG58 cable, keep the cable run as short as possible. Better yet, ditch the RG58 and see if you can replace it with LMR-200 (a high-quality equivalent with half the loss).

What does the picture look like when you're using a *real* low-loss cable? Significantly better, but maybe not as better as you would like. A typical cable for this application--used by at least one 802.11 vendor--is Times Microwave LMR-400. LMR-400 is a very high-quality cable, but it still has a loss of 6.8 dB per 100 feet at 2.4 GHz. This means that, in a 100-foot length of cable, over three quarters of your signal is lost. The moral of the story is clear: keep your access points as close as possible to your antennas. Minimize the length of the transmission line. If you want a roof-mounted antenna, perhaps to cover a courtyard where people frequently have lunch, don't stick your access point in a wiring closet in the basement and run a cable to the roof. If possible, put your access point in a weatherproof enclosure on the roof. If that's not possible, at least put the access point in an attic or crawlspace. There is no substitute for keeping the transmission line as short as possible. Also, keep in mind that transmission lines have a strange ability to shrink when they are routed through walls or conduits. I've never understood why, but even if you measure carefully, you're certain to find that your cable is two feet short. More to the point: the straight-line distance from your access point to the antenna may be only 20 feet, but don't be surprised if it takes a 50-foot cable to make the trip. The cable will probably

have to go around corners and through conduits and all sorts of other misdirections before it arrives at its destination.

If you decide to use an 802.11a product, which operates at 5-GHz, be aware that cable loss will be an even more significant issue. Losses increase with frequency, and coaxial cable isn't particularly effective at 2.4 GHz, let alone 5 GHz.

Finally, there's the matter of antenna connectors. All wireless vendors sell cables in various length with the proper connectors and adapters. I strongly recommend taking the easy way out and buying cables with the connectors preinstalled. Connector failure is one of the most common causes for outages in radio systems, particularly if you don't have a lot of experience installing RF connectors.

Antenna diversity

One common method of minimizing multipath fading is to have *antenna diversity*. Rather than making the antenna larger, radio systems can use multiple antennas and choose the signal from the antenna with better reception. Using multiple antennas does not require sophisticated mathematical theory or signal-processing techniques.

Several wireless LAN vendors have built multiple antennas into wireless network cards. Some vendors even offer the ability to connect multiple external antennas to network cards intended for access points. Antenna diversity is recommended by the 802.11 standard, but it is not required. For environments with large amounts of interference, antenna diversity is a worthwhile option to consider when selecting vendors.

Bring on the heat

Amplifiers are available for increasing your transmitting power. Transmitting amplifiers often incorporate preamplifiers for receiving, helping to improve your weak signal sensitivity. Is an amplifier in your future? It depends. The basic problem is that, as you cover a larger and larger territory, there are more and more stations that can potentially join your network. However, the number of stations that can be

handled by any given access point is fairly limited (see the following section). All in all, more low-power access points provide better service than a smaller number of access points with high-power amplifiers. There may be some applications that are exceptions to the rule (community networks or ISPs in remote areas, for example), but in most situations, high power sounds like a better idea than it really is.

However, if you want to check around and see what's available, SSB Electronics (www.ssbusa.com/wireless.html) and HyperLink Technologies (http://www.hyperlinktech.com/web/amplifiers_2400.html) sell high-power amplifiers for 802.11b service. However, remember:

- To stay within the legal power limit, both for absolute power and ERP.
- That 802.11 is an unlicensed service. If you interfere with another service, it's your problem, by definition. And if a licensed service interferes with you, it's your problem, by definition. Interference is more likely to be a problem if your network covers a large service area and if you are using high power.
- To use equipment that is approved for 802.11 service. Other amplifiers are available that cover the frequency range, but using them is illegal.

> **WARNING:** The FCC does enforce their rules, and their fines are large. If you're in violation of the regulations, they won't be amused, particularly if you're in excess of the power limit or using unapproved equipment.

A word about range

It's tempting to think that you can put up a high-gain antenna and a power amplifier and cover a huge territory, thus economizing on access points and serving a large number of users at once. This isn't a particularly good idea. The larger the area you cover, the more users are in that area--users your access points must serve. Twenty to 30 users for each wireless card in your access points looks like a good upper bound. A single access point covering a large territory may look like a good idea, and it may even work well while the number of users remains small. But if your network

is successful, the number of users will grow quickly, and you'll soon exceed your access point's capacity.

Conducting the Site Survey

When working on the site survey, you must duplicate the actual installation as much as possible. Obstacles between wireless LAN users and access points decrease radio strength, so make an effort to replicate exactly the installation during the site survey.

If access points need to be installed in wiring closets, make sure the door is closed while testing so the survey accounts for the blocking effect of the door on radio waves. Antennas should be installed for the test exactly as they would be installed on a completed network. If office dwellers are part of the user base, make sure that adequate coverage is obtained in offices when the door is closed. Even more important, close any metal blinds, because metal is the most effective radio screen.

Signal measurements should be identical to the expected use of the network users, with one exception. Most site survey tools attempt to determine the signal quality at a single spatial point throughout a sequence of several points in time, and thus it is important to keep the laptop in one location as the measurement is carried out. Taking large numbers of measurements is important because users will move with untethered laptops, and also because the multipath fading effects may lead to pronounced signal quality differences even between nearby locations.

Have several copies of the map to mark signal quality measurements at different tentative access point locations, and note how the antenna must be installed at the location. If multiple antennas were used, note the type and location of each antenna.

Direct-Sequence Channel Layout

Most locations are deploying 802.11 products based on direct-sequence technology because the high-data rate products are based on direct-sequence techniques. Direct sequence underlies both the 2-Mbps DS PHY and the 11-Mbps HR/DSSS PHY. Both standards use identical channels and power transmission requirements.

Direct-sequence products transmit power across a 25-MHz band. Any access points must be separated by five channels to prevent inter-access point interference. Selecting frequencies for wireless LAN operation is based partly on the radio spectrum allocation where the wireless LAN is installed. See Table 15-3.

Table 15-3: Radio channel usage in different regulatory domains

Channel number	Channel frequency (GHz)	US/Canadaa	ETSIb	France
1	2.412	☐	☐	
2	2.417	☐	☐	
3	2.422	☐	☐	
4	2.427	☐	☐	
5	2.432	☐	☐	
6	2.437	☐	☐	
7	2.442	☐	☐	
8	2.447	☐	☐	
9	2.452	☐	☐	
10c	2.457	☐	☐	☐
11	2.462	☐	☐	☐
12	2.467		☐	☐
13	2.472		☐	☐

a 802.11 allows different rules regarding the use of radio spectrum in the U.S. and Canada, but the U.S. Federal Communications Commission and Industry Canada have adopted identical rules.

bNot all of Europe has adopted the recommendations of the European Telecommunications Standards Institute (ETSI). Spain, which does not appear in the table, allows the use of only channels 10 and 11.

cChannel 10 is allowed by all regulatory authorities and is the default channel for most access points when they are initially powered on.

Access points can have overlapping coverage areas with full throughput, provided the radio channels differ by at least five. Only wireless LANs in the U.S., Canada, and Europe that have adopted the ETSI recommendations can operate access points with overlapping coverage areas at full throughput.

After locating the access points, make sure that any access points with overlapping coverage are separated by at least five channels. The cellular-telephone industry uses the "hex pattern" shown in Figure 15-6 to cover large areas.

Figure 15-6. Frequency planning

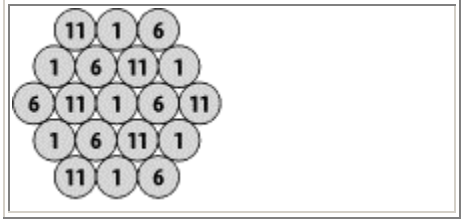

Part of the site survey is to establish the boundaries of access point coverage to prevent more than three access points from mutually overlapping, unless certain areas use multiple channels in a single area for greater throughput.

Limitations of direct-sequence channel layout

One of the problems with 802.11 direct-sequence networks and 802.11b Direct-sequence networks is that there are only three nonoverlapping channels. Four channels are required for nonoverlapping coverage in two dimensions, and more channels are required for three dimensions. When laying out frequency channels in three dimensions, always keep in mind that radio signals may penetrate the floor and ceiling.

Application Performance Characterization

As part of the site survey, take some time to work with the "power users" to ensure that the application performance is adequate. Most applications use web frontends and are relatively tolerant with respect to long delays or coverage dropouts because web browsers retry connections. Terminal emulation and other state-oriented client/server applications may be less tolerant of poor coverage. Part of the engineering in installing a wireless LAN is to tailor the areas of overlapping coverage to offer denser coverage when the applications are less tolerant of momentary drops.

The End of the Site Survey: The Report

After the completion of the site survey, the technical details must be provided to installers to complete the network build-out. Consultants may use the site survey in different ways. Some consultants charge for the site survey and allow the customer or a third party to finish the installation. Value-added resellers may take the same approach or use the site survey to put together an installation bid for the customer.

Depending on the customer's requirements, some or all of the following details may be included in a site survey report:

1. A summary of the requirements from the initial preliminary work.
2. Estimated coverage areas based on the site survey measurements. This may be divided into areas with good coverage, marginal coverage, and weak coverage. It may also site potential trouble spots if the signal strength measurements allow for it.
3. A description of the locations of all access points, along with their configuration. Some elements of this configuration are the following:
    - The access point name
    - Its operating channel
    - Approximate coverage area
    - IP configuration
    - Antenna type and configuration (including direction for directional antennas)
    - Any other vendor-specific information
4. If the customer supplied detailed floor plans or physical network maps, those maps can be returned with detailed access point placement information. Estimated coverage areas can be noted on the map and serve as the basis for frequency reuse planning. Any antenna requirements (external antennas, antenna types, and adjustments to default transmission power) for achieving the noted coverage area should be recorded as well.
5. Many customers appreciate an estimate of the work necessary to install drivers onto any affected laptops. The scope of this item depends a great deal on the sophistication of the management tools used by the customer. For

many, it will be sufficient merely to include a copy of the driver installation instructions as an appendix to the report. Some clients may require low-level details on the driver installation so that the driver installation can be completely automated down to any necessary registry changes on Windows systems.

Installation and the Final Rollout

After the site survey is finished, there should be enough information available to install a wireless LAN. Actual cabling and physical installation may be contracted out or performed by internal staff.

Recordkeeping

Careful documentation is an important part of any network build-out, but it is especially important for wireless LANs because the network medium is invisible. Finding network components is not always a simple matter of cable tracing! To document a wireless LAN, keep the following list in a safe place with the rest of the network maps:

- The site survey report.
- The annotated building blueprints with access point locations, names, and their associated coverage areas. If possible, the blueprints should also indicate areas of marginal or no coverage.
- A separate list of information about the access points in tabular form, which includes the location, name, channel, IP network information, and any other administrative information.

On the Naming of Access Points

Many institutions have naming policies that may dictate DNS names for wireless LAN access points. Device names should be as descriptive as possible, within reason. Companies that provide network service to other users, such as a "hot spot" provider, may wish to keep information about the detailed location of access points

secret from users to keep the physical location of access points secret. Figure 15-7 illustrates a DNS naming convention in which the secrecy of access point locations is not a particular concern, so each name includes the geographic site location, the building name and floor, and the access point number and location on that floor. It also includes the wireless LAN name (SSID).

Figure 15-7. Wireless LAN naming convention for access points

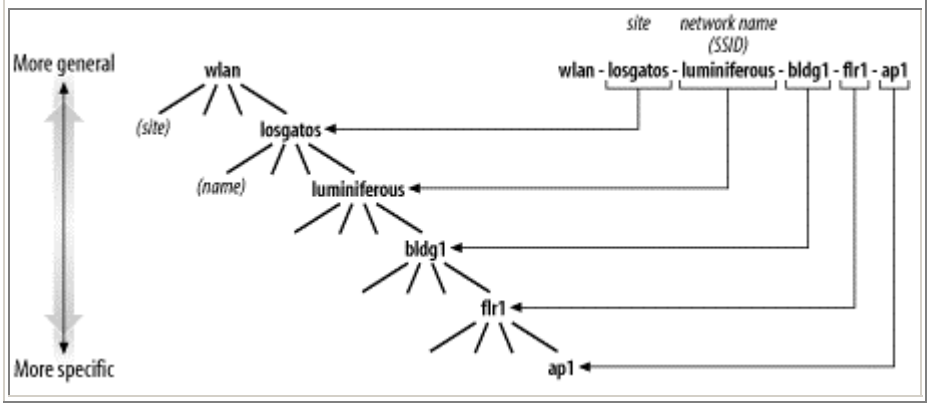

All names are prefixed with *wlan-* to indicate clearly that they are associated with an 802.11 network. The next level of hierarchy is the geographic location of the wireless network; it can often be derived from existing site codes in large companies. Each site may compose a campus and have several buildings, but a single extended service area may offer coverage throughout the entire site at anything up through even midsized installations. The next level of hierarchy is the building identifier, which is often clear from existing conventions. Within a building, the floor number and the location of an access point within the floor can be used to further identify an access point. Figure 15-8 shows how the DNS name can be structured for an access point on the luminiferous network on the second floor of building one at a site in Los Gatos, California (*wlan-losgatos-luminiferous-bldg1-flr2-ap1*). In very large buildings, the access point number might even be replaced by a description of the location of the access point on the floor, such as *-ap-nw-4* for the fourth access point in the northwest corner.

Figure 15-8. Convention for naming wireless LAN stations

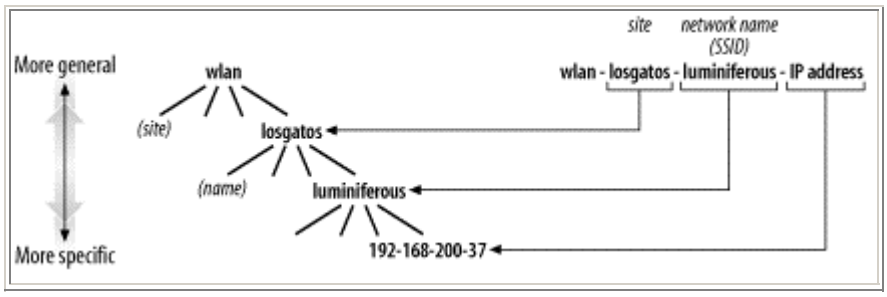

To make troubleshooting easier, follow a convention for the naming of wireless LAN stations. A station on the same network as the access point described previously might be *wlan-losgatos-luminiferous-192-168-200-37*, in which the last set of numbers is the IP address.

Security

After installation is complete, you should execute a second test solely for security. Configuring access points can be time-consuming, detail-oriented work, and it is possible to forget to set a software option here and there. Check the following items to be sure your network is as secure as possible:

- If desired, WEP is enabled on all access points to prevent unauthorized association with the network.
- Lists of MAC addresses allowed to associate with the network have been distributed to each access point.
- Any access controllers in place are properly configured to block initial connections, and they can reach authentication servers to allow user access.
- Any VPN software is properly configured to accept connections from associated stations.
- Access points enforce restrictions on stations allowed to connect for management, and passwords are set.

Some vendors are advertising products that support 802.1x security, particularly in high-end access points. 802.1x requires that you set up a RADIUS authentication server, but the additional security is well worth the effort. On the other hand, 1x products are only just appearing as this book goes to press. Because 802.1x is a

standard, products should interoperate, in theory--but we don't yet know whether they interoperate in practice. If you're not comfortable buying all your equipment from one vendor, you may want to stay away from 802.1x for a while. It's hard to make predictions that mean anything, but here's a guess: given the amount of attention that security has received lately, I expect that the 802.1x situation will be stabilized by the middle of 2002.

Several vendors have proprietary security solutions to replace or supplement WEP--some appear to be preliminary versions of 802.1x. I do not recommend locking yourself into a proprietary solution when a standard solution is available, or nearly available.

---

1. The exception to this general rule is, of course, a network in which Mobile IP is deployed. I am enthusiastic about Mobile IP, especially on wireless networks, but it is far from ubiquitous as I write this book. Most network engineers are, therefore, designing networks without the benefit of network-layer mobility.

2. My definitions here are not meant to be formal. In this section, I'm trying to take a fundamental approach to security by showing how wireless LAN security fails and how some of the failures can be solved by applying solutions the industry has already developed.

3. Or, in the words of one reviewer, "WEP is trash that just gets in the way."

4. The full document title is "Snake Oil Warning Signs: Encryption Software to Avoid." Get your very own copy from http://www.interhack.net/people/cmcurtin/snake-oil-faq.html, among many other places. Bruce Schneier wrote a nice summary of the Snake Oil FAQ in the February 15, 1999, Crypto-Gram newsletter, available from http://www.counterpane.com/crypto-gram-9902.html.

5. NAT devices can block logon traffic and the inter-domain controller chatter used by NT-based networks. See articles Q172227 and Q186340 in the Microsoft Knowledge Base.

6. If you want one more stake, the radio telescope at Arecibo has a gain in excess of 80 dBi.

Back to: 802.11 Wireless Networks: The Definitive Guide