CHAPTER 5

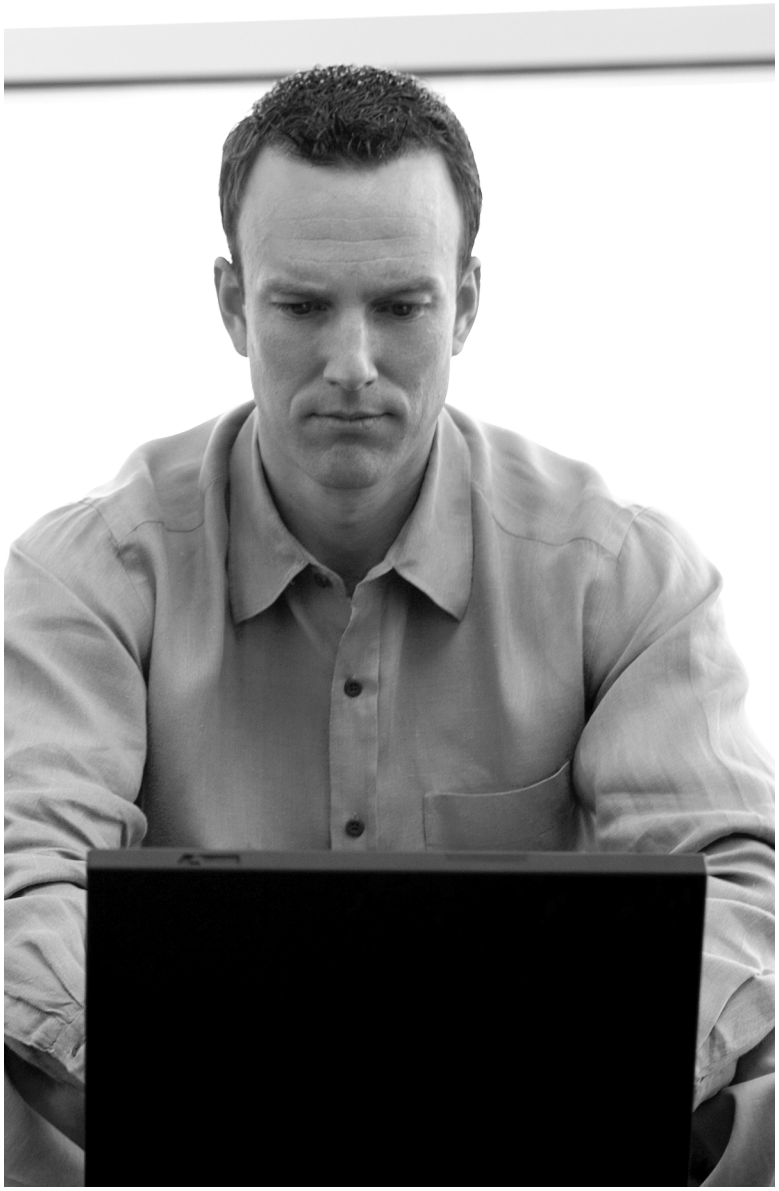

GUIDELINES FOR A SUCCESSFUL ARCHITECTURE AND DESIGN

Part I of this book introduced WLAN technology and familiarized you with its key technical aspects. You learned about the different types of business considerations you need to make to identify, qualify, and quantify the value that WLANs can bring to your organization. You also learned about recommended strategies and practices when initiating the PPDIOO lifecycle of your WLAN. Planning and preparation focused on providing a structured approach for your deployment and highlighted areas that require preparatory work because you need to identify management and technical dependencies that are unique to your circumstances.

As you move through the various lifecycle stages, your focus shifts from strategic to tactical matters. Chapter 2, "Business Considerations," and Chapter 3, "Preparation and Planning" focus on the strategic aspects of setting up your WLAN. Part II of the book covers the next phases of the PPDIOO lifecycle. You learn about architecture, design, implementation, and operations relating to your WLAN.

The difference between architecture and design can be rather vague; however, as a rule of thumb, consider the difference as similar to that between strategic and tactical matters. In both cases, the former is concerned with where to go, whereas the latter is focused on how to actually get there. This chapter covers the strategic aspects of defining your WLAN architecture and takes a look at the tactical design considerations that are specific to WLANs.

This chapter introduces the notion of architecture and provides recommendations for developing a holistic framework that can guide the engineering effort of designing, implementing, and operating the WLAN. You learn about the key components of an effective architecture and identify the balance that must be struck between detail, complexity, and usefulness.

The WLAN design provides the necessary detail on how the solution must be built, integrated, and configured. Because many of the design considerations are identical for wired and wireless networks, this chapter focuses on those considerations that are unique to WLANs. These include the ratio of users to access points, also known as the client-to-AP ratio, the impact of roaming from cell to cell, and the physical placement of the access points.

Finally, this chapter highlights the environmental considerations that are essential for defining a WLAN architecture and design. You learn details about the impact of the physical environment, nearby radio signals, and local governmental regulations and explore the recommended practices for managing these challenges.

Architectural Considerations

Architecture is a framework of components, concepts, and practices that acts as a guide for an underlying design. A robust architecture ensures that the actual WLAN solution meets the predetermined goal for the organization while providing sufficient flexibility to manage the various engineering and operational tradeoffs that WLAN technology requires. As such, it is important that the architecture act only as a guide or baseline and not as a blueprint. This section provides recommendations for setting realistic expectations and guidelines when defining your WLAN's architecture.

WLAN Expectations

The definition of your WLAN architecture should begin with identifying and scoping your expectations and goals. A clear understanding of the business needs for WLANs will simplify alignment between technology solutions and business requirements, and will facilitate the definition of a relevant and specific WLAN architecture. The successful WLAN architecture, therefore, relies on the business considerations, as discussed in Chapter 2, and the provisioning strategy, as outlined in Chapter 3.

When you define your WLAN architecture, focus on two distinct technology alignment challenges:

- Alignment with business requirements
- Alignment with user requirements

To support the business, the WLAN architecture should facilitate and support the generation of a net positive value in the form of strategic, operational, or technological benefits. To effectively support the user, the architecture needs to take into account parameters such as usability, convenience, access, availability, and support. If the WLAN is not easy to use, is subject to poor coverage or uptime, or has little user support, the total WLAN experience will not be positive, resulting in little or no use of the infrastructure investment.

Key Components for an Effective WLAN Architecture

WLANs are justified by benefits such as providing mobility to the workforce, reducing the cost of infrastructure for sporadically used locations, and increasing productivity by keeping mobile users connected. As highlighted in Chapter 2, the business case for WLANs relies on your ability to identify the organizational benefits that WLANs can enable. Identifying which services and applications the WLAN must support is key to building a robust, relevant, and sustainable architecture.

Without a thorough understanding of what is demanded from the wireless communications infrastructure, there is a high probability that you will either undershoot or overshoot supply. Your WLAN architecture thus becomes the vehicle that guides and ensures proper alignment between infrastructure demand and supply. The key components for the step-by-step development of a successful WLAN architecture are

- Determining the goal of the WLAN
- Defining the scope of your WLAN
- Developing your timeframe to deploy
- Considering IT security requirements
- Identifying the types of users and devices you want to support
- Establishing an operational support structure and process

The following sections describe each of these considerations in more detail.

Determining the Goal of the WLAN

Because of increased adoption, more applications and services are being layered onto the WLAN. However, the number of applications utilizing wireless transport is not the only factor that is changing. The characteristics of the applications themselves are changing as well.

Traditionally, WLANs in enterprises were intended only for data traffic. The key applications were typical business productivity tools such as e-mail, web browsers, calendaring tools, and messaging. These applications produce network traffic that is irregular and noncontinuous. Periods with high network utilization are followed by periods of low network utilization, and the duration of both these periods is unpredictable. The applications are considered "bursty" as they load the network in bursts.

As WLANs became more prevalent, they started to become the preferred means of network connectivity. This resulted in bandwidth-intensive and potentially latency-sensitive applications such as video also migrating onto the wireless medium. The challenge created by these applications is that they demanded a different type of service by the network. Best-effort service became insufficient as these applications required high throughout and deterministic behavior.

A common issue with networked applications is that they are developed with little or no consideration for the resources they require from the communications infrastructure. Application developers take into consideration the notion of the network but typically fail to consider bandwidth and latency implications. The (false) assumption is that the network is always available, that bandwidth is unlimited, and that congestion and delays do not occur. As such, even though the applications and the network are tightly coupled, they are typically developed and deployed as independent components.

It is exactly this decoupling that creates the burden of carefully planning your WLAN if you want to successfully support the extension of your applications to the wireless environment. Hence, you should start with the premise that the average application is not aware of the transport medium it is using. They treat the network—wired or wireless—identically.

The challenge of applications not being aware of the network is compounded with WLANs. Indeed, most applications are developed for the wired environment. Specific characteristics of WLANs are their lower throughput and higher latency than their wired equivalents. This is typically not a problem for the bursty applications. However, WLANs can cause additional challenges for applications that demand high data rates or deterministic behavior.

The interaction between applications and the network is only one of the challenges that must be tackled when defining a WLAN architecture. Defining a wireless architecture to support voice and video also introduces specific problems that must be considered. The considerations include provisioning sufficient bandwidth for latency-sensitive voice and video streams, implementing a quality of service (QoS) solution, and ensuring fast roaming capabilities between cells. Refer to Chapter 4, "Supplementary and Complementary Services," for additional details on supporting voice and video in WLAN environments.

Defining the Scope of Your WLAN

The scope or footprint of your WLAN deployment is one characteristic that you can easily define from the start. However you define it—small, limited, partial, full-scale, or ubiquitous—there is a boundary to which you can adhere. Although the scope of your WLAN deployment has a larger impact on the planning and implementation phases, it also plays a role in the architecture.

The architecture must formalize and document the coverage your WLAN provides. The formalization of the scope serves as a guide to ensure that you neither underengineer nor overengineer your WLAN solution. Underengineering occurs when you provide insufficient resources to provide the intended degree of service. Examples include inadequate coverage due to not deploying enough access points or failing to incorporate the proper IT security standards for your organization.

Overengineering is the inverse case. This happens when more resources are supplied than are strictly needed to implement your desired solution. In this scenario, you will have squandered both time and money. An example of overengineering is deploying too many access points. In this case, you could either be overlapping coverage of access points or providing coverage in areas where there is no need for WLANs.

A key consideration when determining the scope of your WLAN is how you intend to provide operational support. Today, enterprises that extend their global reach must deal with an increased number of operational issues. Examples include selecting a scalable strategy and platform for managing the WLANs RF spectrum as well as potentially hundreds of access points and thousands of client devices.

Leverage the scope as defined in the WLAN architecture as a planning tool. This structured approach makes it easier to determine how you offer support at the different levels of the fault resolution path, and how you plan to handle onsite resources for troubleshooting. Refer to Chapter 8, "Management Strategies for Wireless LANs," for operational considerations and recommendations.

Developing Your Timeframe to Deploy

The next step is to develop a timeframe for deployment. In this case, time refers not only to the time when you begin a deployment, but also to the time it takes to complete the deployment. The proper mix of time, as it relates to the

preparation, planning, design and implementation, is fundamental. Two aspects of time concern WLAN architecture and design:

- Making timely decisions
- Implementing the decisions before they outdate themselves

Wireless data networking is no longer a budding technology. It has built up momentum to a point where aspects of the technology are quickly superseded by more advanced features and functions. Figure 5-1 illustrates that as the time to deploy becomes extended, the probability that technology features will make a significant jump becomes greater.

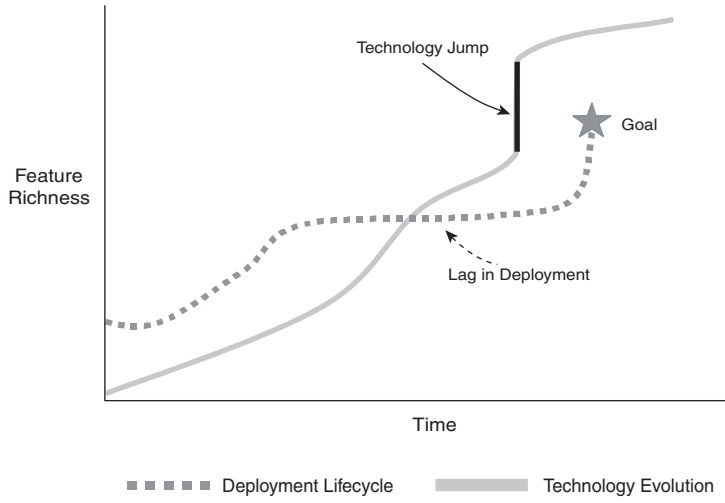

Figure 5-1 *Deployment Versus Technology Evolution*

To manage the time it takes to deploy, adopt the following practices when defining the architecture:

- **Stay familiar with developments in WLAN technology**—Set a goal of staying abreast of standards to ensure that the technology does not date itself. Establish a frequency and process that evaluates the market, and build a matrix to align the technology with the overall business direction.
- **Break up requirements into sections**—By segmenting your business needs, you can align your technology solutions more easily and efficiently. Build a roadmap of the follow-on technologies that can be adopted without requiring a major change in the architecture.

- **Remind yourself that the architecture is only a framework**—The architecture should not become so detailed that it impedes the growth of the network. Maintain standards but avoid defining specific engineering details in your WLAN architecture.

Considering IT Security Requirements

Many enterprises were reluctant to adopt wireless LANs because of the perceived security concerns. The main causes for worry lay in the unbound and uncontrolled nature of the RF transport medium. However, many tools and solutions have been developed that allow you to build a WLAN that is at least as secure as its wired counterpart.

The WLAN architecture plays a key role in securing your WLAN because it explicitly identifies which components must be incorporated as well as their interdependent relationships. The architecture thus effectively defines the security chain, the policies that must be adhered to, and the procedures that must be followed to secure your WLAN.

Because each WLAN component contributes in either a constructive or destructive way to the robustness of the security solution, you must first identify each of these components. Examples of the components include the following:

- Passwords
- Authentication and access methodology
- Encryption and hashing standards
- Devices and their respective operating systems

NOTE Robust passwords form the foundation of security because they are used to unlock the gate to the system. They should be sufficiently strong to prevent easy guessing or hacking. Exhaustive, brute-force methods can uncover all but one-time passwords, therefore, the strategy for common passwords is to make discovery as challenging as possible. Require the use of both uppercase and lowercase alphanumeric characters in addition to special characters. Furthermore, the more characters you use in a password, the stronger it becomes. You should use no less than a 10-character password. Two examples of robust passwords are Ci$cOPr3## (Cisco Press) and G@W1re!3zz (Go Wireless).

Next, the WLAN architecture must define how all the components are interrelated. This process not only ensures that there are no gaps in the security chain, but also that weaker links can be strengthened or more actively managed to provide a holistic and robust security posture. Clearly define the authentication and access methodology, the selected encryption standard, and key management policies.

The security policy and procedures that you define in your WLAN architecture must then be applied in the design, implementation, and operation stages of your WLAN's lifecycle to ensure that security considerations are not only included but also overarching. Chapter 7, "Security and Wireless LANs," details solutions and recommendations for tackling the challenge of constructing a secure WLAN.

Identifying the Types of Users and Devices You Want to Support

An important architectural consideration is the target audience. This aspect of the WLAN architecture is directly related to the people or devices and how they use the WLAN. Every usage profile has its own distinct set of considerations that needs to be included in the architecture. To simplify the challenge of incorporating user and device considerations in your WLAN architecture, start by segmenting the WLAN user-community by identifying common usage profiles.

The concept of user classes was introduced in Chapter 3. Classifying users allows you to determine the degree of relevance of WLANs for subsets of the user community. Classification is performed by grouping users who share common attributes. These profiles are based on the users' characteristic requirements that include their primary applications, degree of mobility, bandwidth and latency restrictions, level of security, and typical hours of operation.

Chapter 3 uses the segmentation in function of mobility needs to define the different user classes. The classes are named standard, mobile, roaming, hot-desk, and guest. You can opt to use the aforementioned classes, or you could, for example, simplify classification into three classes:

- Highly mobile
- Partially mobile
- Nonmobile

Your WLAN architecture must not only identify the different user classes, but also specifically formalize how the WLAN will support each of the respective classes. Figure 5-2 shows an example of a breakout of users in function of mobility needs. A sample of job roles has been added for illustration. Note that this is an illustration only and by no means definitive. For example, a factory worker might need no mobility in one type of role (manufacturing) but require high mobility in another role (warehouse management).

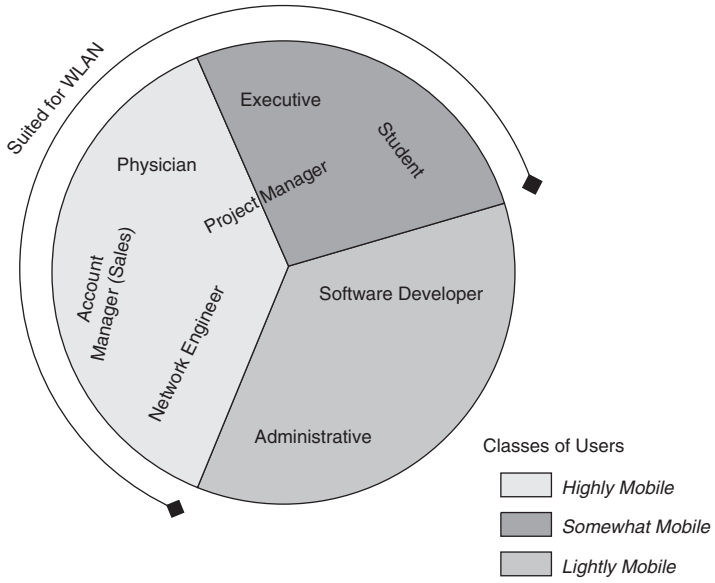

Figure 5-2 *Users by Type and Class*

In addition to the WLAN user considerations, the architecture needs to identify which devices can or must be supported. When thinking about devices, focus on physical attributes. These include tethering, battery life, interoperability, computational horsepower, durability, and control on placement.

For example, clients might be intelligent and mobile such as laptop computers, or dumb and fixed as, for instance, printers and cameras. Handheld scanners can be used for a finite amount of time before their batteries run out. Finally, the placement of PDAs is hard to control, creating potential security hazards.

Different devices have capabilities and limitations that, if supported correctly, ensure performance at the desired expectations. The architecture must frame which devices will be supported by the WLAN and provide guidelines regarding their expected performance, potential pitfalls that are unique to the wireless environment, and problem mitigation strategies.

Plan for the use of devices of different manufacture. Although standards exist, each device is certain to carry its own inherent features, which can result in future compatibility challenges. In the enterprise, where enforcement of standards can be more readily managed, it might be easy to control such issues, however, there are instances where this is not the case. A prime example is a university deployment. In providing access to its student body, a university needs to support a broad assortment of end devices, operating systems, and client software.

Have your WLAN architecture provide a framework and guideline for how you will support your heterogeneous client base in a comprehensive and structured manner. Explicitly define the different user classes and their respective application characteristics and include details on how the specific devices will be supported.

Establishing an Operational Support Structure and Process

You manage the WLAN through provisioning and operational control. Provisioning is the management of systems or devices by sending configurations down to them. You can do this actively or passively. Active provisioning is a repeat process, such as a nightly download. Passive provisioning is performed only when a change is required.

Operational control is the traditional network management whereby systems and devices are monitored. Reactive operational control is performed by waiting for an event or alert and then acting upon it. Proactive operational control is performed by examining data for trends or capturing signs of trouble before an event occurs.

Given the importance of post-deployment support, your WLAN architecture should explicitly define expectations and baseline standards for the WLAN's operational support structure. This should be the case irrespective of the actual management strategy you already have or want to put in place. Strike a careful balance between customization, complexity, and value because too much customization will likely yield diminishing returns. Refer to Chapter 8 for more details on recommended operational practices.

Design Considerations

The previous section provided guidelines for defining the overarching architecture for your WLAN. The framework formalizes the goal, scope, supported device types, and lifecycle management strategy for your WLAN. More specifically, the architecture defines the strategy for the WLAN's security posture and practices, as well as the WLAN's implementation and operational support structure. The architecture does not, however, address detailed design considerations.

The WLAN design provides the necessary detail on how the solution must be built, integrated, and configured. As such, the design of your WLAN specify specify network topologies, how many access points you need to deploy, their make and model, specific AP configurations, where and how you will connect the WLAN to the rest of the network, IP addressing schemes, QoS parameters, access point management passwords, and so on. In short, the design is focused on the physical layout and configuration of the WLAN.

Many of the decisions that must be made during the design of wired networks are directly applicable in the wireless environment. However, there are also distinct considerations that are unique to WLANs, including the following:

- The ratio of users to access points, also known as the client-to-AP ratio
- The impact of roaming from cell to cell
- The physical placement of the access points

This section focuses on the design decisions that need to be made regarding the client-to-AP ratio and roaming capabilities. Chapter 6, "Wireless LAN Deployment Considerations," provides guidelines for identifying the appropriate physical placement of the access points during the implementation of the WLAN.

Client-to-AP Ratio

Many different factors impact the performance of your WLAN. Internal aspects include the shared nature of the communication medium, the access mechanism for the medium, the use of a limited number of communications channels, and the available bandwidth. External factors consist of the number of users, the types of devices communicating across the WLAN, the types of applications used on the network and the degree of mobility that is demanded by the user community.

As outlined earlier in the section "Identifying the Types of Users and Devices You Want to Support," knowing the traffic types and usage patterns on the WLAN is fundamental to designing a solution that not only performs correctly, but also delivers a relatively consistent level of service. As such, providing the WLAN with the proper number of access points is probably the single most contributing factor to creating a WLAN that meets a performance baseline.

The industry has converged on the metric "client-to-access point ratio" to denote the number of users a single access point can consistently support; however, do not take the term "client" at face value. Indeed, a student that uses the WLAN primarily for e-mail and web browsing will have different bandwidth requirements than an engineer using the WLAN mainly for streaming video and computer-aided design (CAD) applications. As such, carefully consider the types of clients and their respective network needs.

NOTE The client-to-AP ratio is expressed as a number such as 10:1. In this case, the number 10 represents the recommended maximum number of clients that can be associated to an AP at any given time. Exceeding this ratio will degrade the expected performance.

Three different strategies can be used to determine what the correct client-to-AP ratio is for your environment. You can perform benchmark tests to identify exactly what works, you can classify users and traffic types as in Table 5-1 to generate more granular client-to-AP ratio specifications, or you can simply adopt client-to-AP ratio guidelines that have been published by most vendors. Each strategy has its merits and drawbacks.

Benchmarking enables the most precise identification of the client-to-AP ratio. Local variations are measured and the ratio can be optimized depending on the exact user profiles and needs. However, not only is this approach time and resource intensive, but it also creates a dated snapshot. If the environment changes, for example, and the HR and engineering departments introduce new software with different traffic signatures, the benchmarks will no longer be accurate.

By classifying both traffic and users, as detailed in Chapter 3, some degree of customization can be captured. The process is relatively straightforward and can

be performed by your network architects and designers. A challenge that you will likely face with this method is the identification of the correct segmentation of the users and traffic types. Don't reinvent the wheel. Follow the classification guidelines as set forth in your architecture. Given the benefits of more accurately identifying a client-to-AP ratio that yields a more consistent and satisfactory WLAN user experience, we recommend that you adopt this approach.

The final strategy is to accept the recommended client-to-AP ratio as published by the WLAN equipment vendor. Even though this is the easiest solution, there is potential for over- or underprovisioning the number of access points because the information provided by the vendor does not consider your specific user-base requirements. However, use the WLAN vendor's published recommendations as a sanity check.

Roaming

Roaming occurs when a device moves its association from one access point to another. By moving the association, the device has effectively traversed the basic service set (BSS) boundary and moved into a new one. However, roaming is not limited to crossing BSS boundaries.

As mentioned in Chapter 1, "Introduction to Wireless LAN Technologies," the BSS is equivalent to a Layer 2 network. Multiple BSSs can be grouped together into an extended service set (ESS), which equates to a Layer 3 network. As such, changing the association from one access point to another can not only cause the client to roam across BSS boundaries, but also ESS boundaries.

Authentication is not the only area that is affected when a user moves its association from one access point to another. Roaming across BSS boundaries creates the following three challenges:

- Authentication
- Performance
- ESS boundaries

Each vendor offers its own solution for these challenges, and each solution has its own strengths and weaknesses. In the end, it is important to understand the impact of roaming. The following sections take a closer look at the challenges that are created by roaming and provide recommendations for addressing them.

Authentication

If you opt to use authentication to secure your WLAN, switching association from one AP to another triggers a re-authentication process. The new AP does not know that the client is permitted to associate and, therefore, the client must go through the entire authentication process. As the number of times a station roams and the number of stations roaming increases, latency can be introduced due to the authentication traffic and the authentication processing overhead that is handled by the AP.

Note that authentication does not occur only when a client roams. To increase the robustness of WLAN security, it is not uncommon that authenticated credentials expire after a certain amount of time. When this occurs, the station is forced to re-authenticate. In this scenario, a station authenticates multiple times over the duration of its association with the same access point even though it is not physically roaming.

Some WLAN products provide methods to reduce the number of authentication requests that are sent to the authentication, authorization, and accounting (AAA) infrastructure. This process is often known as *fast roaming*, because the authenticated status of the client is stored locally in the access point or controller, thereby avoiding the need to contact the back-end AAA server directly. This reduces the time for authentication (hence "fast roaming") and the load on the AAA servers themselves.

Performance

Performance is not limited to the throughput that a client can achieve. It is also directly related to the client keeping its network connection and communication session intact. When roaming, there is a small amount of time during either authentication or association during which the client will effectively be without a link. The duration of the lost link will determine if and how applications will be impacted. Note that last roaming was specifically conceived to make this link loss during authentication almost unnoticeable to end users.

Applications exhibit a distinctive sensitivity to the duration of a lost link. Transactional applications such as e-mail and web browsing are relatively insensitive, whereas real-time applications such as voice and video are highly sensitive. Ensure that you enable fast roaming to make authentication occur promptly enough to not affect the core WLAN application suite.

ESS Boundaries

As mentioned earlier, roaming occurs when a station moves its association from one access point to another. This effectively makes the station jump from one BSS cell into the next. As long as the client remains in the same ESS, its IP address remains valid and the Layer 3 session can be maintained.

If, however, the station crosses an ESS boundary, it effectively moves into a different Layer 3 network. The IP address that was assigned for the old ESS is invalid, and all active IP sessions terminate as traffic directed toward the station is incorrectly routed. To remediate this routing problem, the client must release its old IP address and request a new one for the subnet that it now finds itself in.

To keep the IP sessions alive, some mechanism is needed to transfer the active connections. A method of achieving this is by empolying Mobile IP, which is an open protocol that comes in different forms but allows clients to move between Layer 3 networks or subnets. However, keep in mind that Mobile IP is no longer the primary mobility method for most vendors. Because it requires client software, it is currently used only in "extreme" roaming situations like those found in moving vehicles with multiple available network types. Most vendors today use some kind of tunneling technology to hide the fact that the user has crossed a Layer 3 network boundary. This tunneling solution is similar to that used for remote VPN access. In essence, a logical overlay of multiple ESSs is instantiated by means of the tunnels, thus enabling roaming without Layer 3 hazards.

If you do not opt to implement solutions that provide Layer 3 roaming capabilities, carefully plan the layout of your WLAN subnets to address this challenge. Avoid creating multiple ESSs in areas where users typically roam. For example, because users typically move around on a floor, create a single ESS per floor. However, a floor-by-floor model can have problems in certain buildings where there is strong signal propagation between floors. In these types of buildings, users can accidentally roam between floors, creating the problems previously described. Carefully measure signal strength on each floor and fine-tune the radio's signal power to avoid it propagating between floors.

Also, consider recommended practices for sizing IP subnets. Subnets that are too large can experience performance issues because of excessive IP broadcast traffic. Adopt the recommended IP addressing practices when designing your WLAN. Plan carefully and strike a balance.

Environmental Considerations

The environment—be it a building, country, or climate—in which the WLAN operates plays a critical role in defining the architecture and design of WLANs. Chapter 1 introduced the various environmental factors that have an impact on the performance of the WLAN. Examples included the attenuation and distortion of radio signals by various materials and the multipath effect.

The architecture should account for the variables of the environment without actually providing specific details on remediation methods. The design, however, must include specifications on how the WLAN will accommodate local variations.

When defining your WLANs architecture and design, you need to consider the following three environmental matters:

- Physical attributes of the surroundings
- RF environment
- Local governmental regulations

The following sections describe each point in detail.

Physical Attributes of the Surroundings

Physical attributes of the surroundings include both the placement and type of obstacles in open spaces. The obstacles are any foreign objects that the propagating RF signals encounter. The objects can be static, such as furniture, walls, warehouses, and buildings, or dynamic like cars, forklifts, and even people.

The exact type or movement of the obstacle is not important. What is important is that all matter affects the radio signal by modifying the profile and strength of the original signal, which is known as distortion and attenuation, respectively. As the signal's quality and strength decreases, the receiving stations have a harder time reconstructing the original message, and this is accompanied by a reduction in throughput.

Include considerations about the physical environment in your architecture by defining what the minimum acceptable signal quality and strength should be because they determine which throughput your WLAN can sustain. For example, requirements for a signal that must consistently support 54 Mbps throughput will be a lot more stringent than an environment where throughput can throttle between

54 Mbps and 11 Mbps. These guidelines should then be used for the design of the WLAN to determine how many access points are required and where they should be placed.

RF Environment

Physical obstacles are not the only kind of entities that can impact the strength and quality of an RF signal. Other RF signals that are in the vicinity interfere with the original signal and modify its profile.

Whereas the visible concerns can be managed in a straightforward manner, the invisible cannot. Even in controlled deployments, you can expect to contend with other nearby WLAN deployments. Furthermore, devices like wireless phones, microwaves, handheld radios, and Bluetooth devices will have some impact on RF signal quality because (in most cases) they share the same RF space.

The best way to combat the challenge of interference is to carefully and purposely design your WLAN cells. Fix the throughput rate of your cells. Building your WLAN with well-defined cells aids in the control and troubleshooting of these unknowns. An additional benefit of carefully controlling the footprint of the radio cells is a higher degree of security. Chapter 7 covers the security considerations in more detail.

WLAN protocols are designed to throttle throughput in function of the strength and quality of the signal. As the footprint of the cell is related to the throughput, varying rates result in changing cell sizes. For example, in 802.11b, cells that are fixed at 11 Mbps are significantly smaller than those that are fixed at 2 Mbps. Pegging the throughput rate creates a fixed cell-size that is an easier to use building block for designing your WLAN. The ability to design the network with standard and well-known parameters also makes it easier to set the expectations of the user and troubleshoot the WLAN.

Local Governmental Regulations

One of the most overlooked aspects of WLANs in the global enterprise is the different regulations and limitations that are imposed on WLANs and, specifically, the RF spectrum in which they operate. Regulations surrounding RF are managed

by both national and regional bodies, and significant disparities can exist between respective local regulations. Be aware of potential differences in RF regulations and know which regulatory bodies are relevant in your specific case and where to consult upon them. Include both the regulatory requirements and the correct local solutions in both your architecture and design.

Summary

This chapter discussed the key architectural, design, and environmental considerations that are required for WLANs. It emphasized the need for the architecture to be a framework as opposed to a blueprint, thus providing flexibility for the designer. You learned about guidelines and recommended practices for defining a robust WLAN architecture, including the following:

- Determining the goal of the WLAN
- Defining the scope of your WLAN
- Developing your timeframe to deploy
- Considering security requirements and implications early
- Identifying the types of devices you want to support
- Establishing an operational support structure and process
- Adopting a financially responsible and conservative position
- Confirming the staffing model for building and maintaining the WLAN

This chapter also discussed the most important design considerations that are specific to WLANs. The need and methods for determining the correct client-to-AP ratio were covered as well as the challenges that are created by roaming of stations.

Finally, the environmental considerations that are essential for defining a WLAN architecture and design were highlighted. The impact of the physical environment, nearby radio signals, and local governmental regulations was looked at in addition to recommended practices for managing these challenges.