

Microsoft Solutions for Security

*Securing Wireless LANs with
Certificate Services*

Planning Guide

Release 1.6

Microsoft®

Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e – mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e – mail address, logo, person, place or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2004 Microsoft Corporation. All rights reserved.

Microsoft, Windows, Active Directory, Windows NT, and Windows Server 2003 are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners

Planning Guide Overview

The Planning Guide is intended for IT architects and provides:

- Business and technical reasons for implementing wireless security.
- Strategies for wireless security.
- Detailed discussion of the design decisions in the solution and the solution components.

In addition, the design chapters include extended discussions of technical topics and other background information to help you customize the design if required.

The Planning Guide contains the following chapters:

- Chapter 2: Deciding On a Secure Wireless Networking Strategy
- Chapter 3: Secure Wireless LAN Solution Architecture
- Chapter 4: Designing the Public Key Infrastructure
- Chapter 5: Designing a RADIUS Infrastructure for Wireless LAN Security
- Chapter 6: Designing Wireless LAN Security Using 802.1X

Table of Contents

Chapter 2: Deciding on a Secure Wireless Networking Strategy	1
Introduction	1
Overview of Wireless Solutions.....	1
The Argument for Wireless Networking	4
Benefits of Wireless LANs.....	4
Security Concerns with Wireless LANs	5
How to (Really) Secure Your WLAN	8
Protecting the WLAN with 802.1X Authentication and Data Encryption	8
Other Approaches to WLAN Security.....	13
Selecting the Right WLAN Options	20
Deciding on the Right WLAN Security Solution	20
Choosing Between Dynamic WEP and WPA.....	22
Summary.....	23
References.....	23
Chapter 3: Secure Wireless LAN Solution Architecture	25
Introduction	25
Chapter Prerequisites.....	25
Chapter Overview.....	26
Conceptual Design.....	27
Solution Design Criteria	31
Target Organization.....	31
Organization Requirements.....	33
Solution Design Criteria.....	34
Solution Logical Design	35
Review of Conceptual Design	35
Logical Design.....	35
Extending the Design	41
Design Criteria Re-Evaluated	44
Summary.....	45
Chapter 4: Designing the Public Key Infrastructure	47
Introduction	47
Chapter Prerequisites.....	48
Chapter Overview.....	48
Defining Certificate Requirements	51
Creating a Certificate Practices Statement	51
Identifying Certificate Applications	51
Defining Certificate Clients	53
Defining Certificate Security Requirements	55
Designing the Certification Authority Hierarchy	61
Selecting a Trust Model.....	61
Defining Certification Authority Roles.....	66
Supporting IT Infrastructure.....	81
Configuring CDP and AIA Paths	83
Extending Your CA Infrastructure.....	85
Configuring Certificate Profiles	88
Defining Certificate Parameters	88
Defining Certificate and Key Lifetimes	89
Mapping Security Requirements onto Certificate Parameters.....	91
Mapping Certificate Requirements onto Certificate Template Parameters.....	91
Creating Certificate Templates.....	93
Creating a Certificate Management Plan	94

Selecting Enrollment and Renewal Methods	94
Mapping Certificates to Identities	95
Creating Certificate Policies	96
Defining Certificate Revocation Policy	97
Planning for Key and Data Recovery	100
Summary	101
More Information	101
Chapter 5: Designing a RADIUS Infrastructure for Wireless LAN Security	103
Introduction	103
Chapter Prerequisites	103
Chapter Overview	104
Using IAS for Network Access Management	105
Identifying Your Organizational Network Access Management Requirements	105
Using IAS for Wireless Network Access Management	106
Identifying Prerequisites for the Solution	107
Active Directory Considerations	107
Pre-existing RADIUS Infrastructure	108
Designing the RADIUS Infrastructure	109
Determining the Role of IAS as a RADIUS Server	109
Understanding Server Failover and Load Balancing	110
Establishing Logging Requirements	114
Choosing to Centralize or Distribute Servers	117
Determining the Number and Location of Servers	119
Determining Co-Location of IAS with Other Services	119
Estimating RADIUS Server Load	120
Estimating Server Hardware Requirements	122
Determining Server Software Requirements	124
Creating a Management Plan	125
Change and Configuration Management	125
Planning for Service Recovery	125
Planning Administrative Permissions	125
Security Monitoring and Auditing	128
Summary	130
More Information	130
Chapter 6: Designing Wireless LAN Security Using 802.1X	131
Introduction	131
Chapter Prerequisites	131
Chapter Overview	132
Using 802.1X and Encryption to Secure WLANs	134
Deciding on Certificates or Passwords	135
Solution Prerequisites	137
Client Computer Requirements	137
Required Server Infrastructure	137
Required WLAN Equipment	137
Considering WLAN Security Options	138
Selecting User- and/or Computer-Based Authentication	138
Determining Network Authorization Requirements	140
Choosing a Client Configuration Strategy	142
Determining Traffic Encryption Requirements	142
Choosing a WLAN Migration Strategy	143
Wireless Network Infrastructure Design	144
Wireless Network Group Policy Considerations	146
Determining Software Settings Required for 802.1X WLANs	147
Configuring Remote Access Policies	147

Configuring Connection Request Policies	148
Configuring Group Policy for Client Computers	151
Additional Considerations	156
Supporting Roaming Profiles and Roving Users	156
Supporting Clients Without Wired LAN Connections	156
Summary	157
More Information	157

2

Deciding On a Secure Wireless Networking Strategy

Introduction

Wireless local area network (WLAN) technology is a controversial topic. Organizations that have deployed WLANs are concerned about whether they are secure. Other organizations that have not deployed them are worried about missing out on user productivity benefits and lowering their total cost of ownership (TCO). And confusion still exists about whether a WLAN is safe to use in corporate environments.

Ever since weaknesses in the first generation of WLAN security software were discovered, analysts and network security firms have been striving to resolve them. Some of these efforts have contributed significantly to improved wireless security. But others have had their share of flaws: some introduce a different set of security vulnerabilities; some require costly proprietary hardware; and others avoid the question of WLAN security entirely by layering on another, potentially complex security technology, such as virtual private networks (VPN).

In parallel, the Institute of Electrical and Electronic Engineers (IEEE), along with other standards organizations, have been diligently redefining and improving wireless security standards to enable WLANs to stand up to the hostile security environment of the early twenty-first century. Thanks to the efforts of these organizations and industry leaders, “WLAN security” is no longer an oxymoron. Now you can deploy and use WLANs with a high level of confidence in their security.

This chapter introduces two WLAN security solutions from Microsoft and answers questions about best practices for securing WLANs.

Overview of Wireless Solutions

The main objective of this chapter is to help you decide on the most suitable way to secure a WLAN in your organization. To do this, the document deals with four main areas:

- Addressing security concerns associated with WLANs .
- Using secure WLAN standards
- Adopting alternative strategies such as a virtual private network (VPN) and Internet Protocol security (IPsec)
- Selecting the right WLAN options for your organization.

Microsoft has produced two WLAN solutions, based on open standards from standards organizations such as the IEEE, the Internet Engineering Task Force (IETF), and the Wi-Fi Alliance. The two solutions are titled *Securing Wireless LANs with Certificate Services* and *Securing Wireless LANs with PEAP and Passwords*. As the names suggest, the former uses public key certificates to authenticate users and computers to the WLAN, while the latter uses simple user names and passwords. However, the basic architecture of the two solutions is similar. Both are based on Microsoft® Windows Server™ 2003 infrastructure and client computers running Microsoft Windows® XP and Microsoft Pocket PC 2003.

Although not apparent from the titles, the intended audience is different for each solution. The *Securing Wireless LANs with Certificate Services* solution is primarily aimed at large organizations with relatively complex information technology (IT) environments; *Securing Wireless LANs with PEAP and Passwords* is a significantly simpler solution that much smaller organizations can easily deploy.

However, this is not to say that large organizations cannot use password authentication or that certificate authentication is not suitable for smaller organizations. The use of these technologies simply reflects the type of organization in which they are more likely to be applied. The following figure includes a simple decision tree to help you select the WLAN solution that is the most appropriate for your organization.

The three main technology options available to implement WLAN security are to use:

- Wi-Fi Protected Access (WPA) Pre-shared Key (PSK) for very small businesses and home offices.
- Password-based WLAN security for organizations that do not want to use certificates.
- Certificate-based WLAN security for organizations that want to deploy certificates.

These implementation options are explained later in this chapter, as well as the possibility of merging the features of the last two options to produce a hybrid solution.

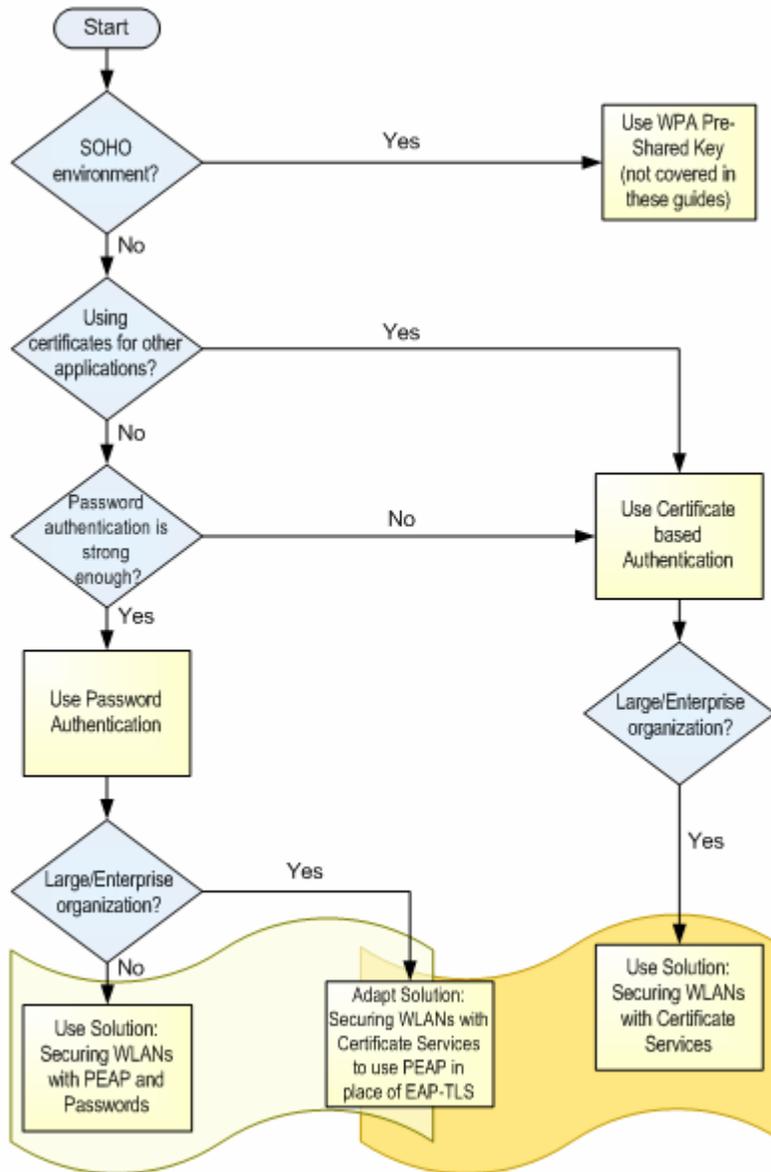


Figure 2.1
The decision tree for the two Microsoft wireless LAN solutions

The Argument for Wireless Networking

It is easy to understand the appeal of WLANs for businesses today. WLAN technology has been around in one form or another for nearly a decade, but it has failed to gain acceptance until relatively recently. Only after reliable, standardized, and low cost technology met the growing desire for more flexible ways of working with more connectivity did WLAN adoption really start to take off. The rapid adoption of this technology, though, has also brought to light a number of serious security weaknesses with first generation WLANs. This section discusses both the pros (the functionality) and cons (security) of WLANs.

The Benefits of Wireless LANs

The benefits of WLAN technology fall into two main categories: core business benefits and operational benefits. Core business benefits include improved employee productivity, quicker and more efficient business processes, and greater potential for creating entirely new business functions. Operational benefits include lower costs of management and lower capital expenditure.

Core Business Benefits

The core business benefits of WLANs arise from the increase in flexibility and mobility of your workforce. People are freed from their desks and can easily move around the office without losing their connections to the network. It is helpful to look at some examples of how increased mobility and network flexibility can benefit businesses.

- Mobile workers moving between offices save time and heartache by staying connected to the corporate local area network (LAN). Users can connect almost instantly from any physical location with wireless coverage, and they do not have to hunt for network ports, cables, or IT staff to help connect them to the network.
- Knowledge workers can stay in touch wherever they are in the building. Using e-mail, electronic calendars, and chat technologies, your staff can remain online even while in meetings or working away from their desks.
- Online information is always available. Meetings no longer are delayed while someone dashes out to retrieve the report of last month's figures or get a presentation update. This can significantly improve the quality and productivity of your meetings.
- Organizational flexibility is also enhanced. As teams and project structures change, desk moves, or even whole office moves are quick and easy because people are no longer wired to their desks.
- Integration of new devices and applications into the corporate IT environment improves significantly. Until recently, devices like personal digital assistants (PDA) and tablet PCs were often executive playthings on the margins of corporate IT; these can become far more integrated and useful when organizations are wireless-enabled. Workers and business processes that were previously untouched by IT can benefit from the provision of wireless computers, devices, and applications into formerly network-free areas, such as manufacturing shop floors, hospital wards, stores, and restaurants.

Different organizations will experience different benefits; which of these are relevant to your organization depends on many factors such as the nature of your business and the size and geographic distribution of the workforce.

Operational Benefits

The main operational benefits of WLAN technology are lower capital and operational costs that can be summarized as follows:

- The cost of provisioning network access to buildings is substantially lowered. Although most office space is cabled for networks, many other workspaces such as factory floors, warehouses, and stores are not. You can also provision networks at locations where wired networks would be impractical, for example, outdoors, at sea, or even in a battlefield.
- You can easily scale the network to respond to different levels of demand as the organization changes, even from day-to-day, if required. For example, it is far easier to deploy a higher concentration of wireless access points (AP) at a given location than to increase the number of wired network ports.
- Capital cost no longer is tied to building infrastructure because you can move your wireless network infrastructure to a new building relatively easily. Wiring is usually a permanent fixture cost.

The Security Concerns with Wireless LANs

Despite the benefits, a number of security concerns with WLANs have limited organizations from adopting them, particularly in security-conscious sectors such as finance and government. Though the risks of broadcasting unprotected network data to anyone in the vicinity might seem evident, a surprising number of WLANs are installed without any security features enabled. The majority of businesses have implemented some form of wireless security. However, the security is usually only in the form of basic, first generation features, which offer inadequate protection by today's standards.

When the first IEEE 802.11 WLAN standards were being written, security was nowhere near as big a concern as it is today. The level and sophistication of threats was much lower and the adoption of wireless technology was in its infancy. It is against this background that the first generation WLAN security scheme, known as Wired Equivalent Privacy (WEP), emerged. WEP underestimated the measures needed to make the security of the air "equivalent" to the security of a wire. In contrast, modern WLAN security methods are designed to work in a hostile environment like the air where there are no clear physical or network perimeters.

It is important to distinguish between first generation static WEP (which uses a shared password to protect the network) and security schemes that use WEP encryption coupled with a strong authentication and encryption key management. The former is a complete security scheme including authentication and data protection and is referred to in this chapter as "*static WEP*." On the other hand, *dynamic WEP* defines only the data encryption and integrity method, which are used in the more secure solutions described later in this chapter.

The security weaknesses discovered in static WEP create vulnerabilities for WLANs protected by it that make them subject to several types of threats. Freely available "audit" tools like Aircrack and WEPCrack make breaking into static WEP-protected wireless networks a trivial task. Unsecured WLANs are obviously exposed to these same threats as well; the difference is that less expertise, time, and resources are required to carry out such attacks on unsecured WLANs.

Before discussing how modern WLANs security solutions work, it is worth reviewing the principal threats to WLANs. These threats are summarized in the following table.

Table 2.1: Main Security Threats for WLANs

Threat	Threat description
Eavesdropping (disclosure of data)	Eavesdropping on network transmissions can result in disclosure of confidential data and unprotected user credentials, and the potential for identity theft. It also allows sophisticated intruders to collect information about your IT environment, which they can use to mount an attack on other systems or data that might not otherwise be vulnerable.
Interception and modification of transmitted data	If an attacker can gain access to the network, he or she can use a rogue computer to intercept and modify network data communicated between two legitimate parties.
Spoofing	Ready access to an internal network allows an intruder to forge apparently legitimate data in ways that would not be possible from outside the network, for example, a spoofed e-mail message. People, including system administrators, tend to trust items that originate internally far more than something that originates outside the corporate network.
Denial of service (DoS)	A determined assailant may trigger a DoS attack in a variety of ways. For example, radio-level signal disruption can be triggered using something as low-tech as a microwave oven. There are more sophisticated attacks that target the low-level wireless protocols themselves, and less sophisticated attacks that target networks by simply flooding the WLAN with random traffic.
Free-loading (or resource theft)	An intruder may want nothing more sinister than to use your network as a free point of access to the Internet. Though not as damaging as some of the other threats, free-loading will at the very least not only lower the available level of service for your legitimate users, but may also introduce viruses and other threats.
Accidental threats	Some features of WLANs make unintentional threats more real. For example, a legitimate visitor may start up a portable computer with no intention of connecting to your network, but then is automatically connected to your WLAN. The visitor's portable computer is now a potential entry point for viruses onto your network. This kind of threat is only a problem in unsecured WLANs.
Rogue WLANs	If your company officially has no WLAN you may still be at threat from unmanaged WLANs springing up on your network. Low priced WLAN hardware bought by enthusiastic employees can open unintended vulnerabilities in your network.

Security concerns with WLANs, focused on static WEP, have received a great deal of attention in the media. Despite the fact that good security solutions exist to combat these threats, organizations of all sizes remain wary of WLANs; many have halted deployment of WLAN technology or even banned it. Some of the key factors contributing to this confusion and the popular misconception that WLANs and nonsecure networks go hand-in-hand include:

- Widespread uncertainty exists over which WLAN technology is secure and which is not. Businesses are suspicious of all WLAN security measures after a succession of flaws were discovered in static WEP. The bewildering list of official standards and proprietary solutions claiming to resolve the problems has done little to clear up the confusion.
- Wireless is invisible; for network security administrators this is not just psychologically unsettling, it poses a real security management problem. Whereas you can actually see an intruder plugging a cable into your wired network, intrusion into WLANs is much less tangible. The traditional physical security defenses of walls and doors that help guard your wired network are no protection from a “wireless” attacker.
- There is now much greater consciousness of the need for information security. Businesses demand much higher levels of security in their systems and do not trust technologies that may create new security vulnerabilities.
- As a corollary to this increasing security awareness, legislative and regulatory requirements that govern data security are appearing in a growing number of countries and industry sectors. One of the best known examples of this is the United States government’s Health Insurance Portability and Accountability Act of 1996 (HIPAA) that governs the handling of personal healthcare data.

How to (Really) Secure Your WLAN

Since the discovery of WLAN security weaknesses, leading network vendors, standards organizations, and analysts have focused a great deal of effort on finding remedies for these vulnerabilities. This has yielded a number of responses to the concerns over WLAN security. The principal alternatives are:

- Not to deploy WLAN technology.
- Stay with 802.11 static WEP security.
- Use a VPN to protect data on the WLAN.
- Use IPsec to protect WLAN traffic.
- Use 802.1X authentication and data encryption to protect the WLAN.

These alternative strategies are listed in order of the least to the most satisfactory based on a combination of security, functionality, and usability, although this is somewhat subjective. Microsoft favors the last of these alternatives: using 802.1X authentication and WLAN encryption. This approach is discussed in the following section and then measured against the list of major WLAN threats identified earlier in Table 2.1. The principal advantages and disadvantages of the other approaches are also discussed later in the chapter.

Protecting the WLAN with 802.1X Authentication and Data Encryption

There are many good reasons to recommend this approach (although its title and array of obscure terminology are not among them). Before discussing the advantages of solutions based on this approach, it is important to clarify some of the terminology and explain how such a solution works.

Understanding WLAN Security

Protecting a WLAN involves three major elements:

- Authenticating the person (or device) connecting to the network so that you have a high degree of confidence that you know who or what is trying to connect to the network.
- Authorizing the person or device to use the WLAN so that you control who has access to the network.
- Protecting the data transmitted on the network so that it is secured from eavesdropping and unauthorized modification.

In addition to these areas, you may require an auditing function for your WLAN, although auditing is primarily a means to check and reinforce the other security elements.

Network Authentication and Authorization

Static WEP security relies on a simple shared secret (password or key) for authenticating users and devices to the WLAN. Anyone possessing this secret key can access the WLAN. Cryptographic flaws in WEP present an opportunity for an attacker to use readily-available tools to discover the static WEP key in use on a WLAN. The original WEP standard also does not provide a method to automatically update or distribute the WEP key, making it extremely difficult to change it. Once a static WEP WLAN is cracked it usually stays cracked.

To provide a much stronger method of authentication and authorization, Microsoft and a number of other vendors proposed a WLAN security framework using the 802.1X protocol. The 802.1X protocol is an IEEE standard for authenticating access to a network and, optionally, for managing keys used to protect traffic. Its use is not limited to wireless networks; it is also implemented in many high-end wired LAN switches.

The 802.1X protocol involves the network user, a network access (or gateway) device such as a wireless AP, and an authentication and authorization service in the form of a RADIUS Remote Authentication Dial-In User Service (RADIUS) server. The RADIUS server performs the job of authenticating the users' credentials and authorizing the users' access to the WLAN.

The 802.1X protocol relies on an IETF protocol called the Extensible Authentication Protocol (EAP) to carry out the authentication exchange between the client and the RADIUS server. This authentication exchange is relayed by the AP. EAP is a general protocol for authentication that supports multiple authentication methods, based on passwords, digital certificates, or other types of credential. extensible

Because EAP provides you with authentication method options, there is no one EAP standard authentication type to be used. Different EAP methods, using different credential types and authentication protocols, may be appropriate for different circumstances. The use of EAP methods in WLAN authentication is discussed in a later section of this chapter.

WLAN Data Protection

Deciding on the 802.1X authentication and network access comprises only part of the solution. The other significant solution component is what you will use to protect the wireless network traffic.

The flaws in WEP data encryption described earlier might have been ameliorated if static WEP had included a method to automatically update the encryption keys regularly. Tools for cracking static WEP need to collect between one and ten million packets encrypted with the same key. Because static WEP keys often remain unchanged for weeks or months, it is usually easy for an attacker to collect this amount of data. As all computers on a WLAN share the same static key, an attacker can collect data transmissions from all computers on the WLAN to help discover the key.

Using a solution based on 802.1X allows you to frequently change the encryption keys. As part of the 802.1X secure authentication process, the EAP method generates an encryption key that is unique to each client. To mitigate attacks against the WEP key, (described earlier), the RADIUS server regularly forces the generation of new encryption keys. This allows you to use WEP encryption algorithms (found in most current WLAN hardware) in a much more secure way.

WPA and 802.11i

Although WEP with 802.1X dynamic rekeying is secure for most practical purposes, there are a few lingering problems including:

- WEP uses a separate static key for global transmissions like broadcast packets. Unlike the per-user keys, the global key is not renewed regularly. Although confidential data is unlikely to be transmitted using broadcast, using a static key for global transmission gives attackers the potential to discover information about the network such as IP addresses, and computer and user names.
- WEP protected network frames have poor integrity protection. Using cryptographic techniques, an attacker can modify information in the WLAN frame and update the frame's integrity check value without the receiver detecting it.

- As WLAN transmission speed improves and computational power and cryptanalytic techniques improve, WEP keys must be renewed with greater frequency. This may place an unacceptable load on the RADIUS servers.

To address these problems, the IEEE is working on a new WLAN security standard called 802.11i; also known as Robust Security Network (RSN). The Wi-Fi Alliance, a consortium of the leading Wi-Fi vendors, has taken what is essentially an early release of 802.11i and published it in an industry standard known as WPA (Wi-Fi Protected Access). WPA includes a large subset of features of 802.11i. By publishing WPA, the Wi-Fi Alliance has been able to mandate adherence to WPA for all equipment bearing the Wi-Fi logo, and allowed Wi-Fi network hardware vendors to offer a standardized high security option in advance of the publication of 802.11i. WPA brings together a set of security features that are widely regarded as the most secure techniques currently available for securing WLANs.

WPA includes two modes; one using 802.1X and RADIUS authentication (simply known as WPA), and another simpler scheme for SOHO environments using a pre-shared key (known as WPA PSK). WPA couples robust encryption with the strong authentication and authorization mechanism of the 802.1X protocol. And WPA data protection eliminates the known vulnerabilities of WEP by providing the following:

- A unique encryption key for each packet.
- A much longer initialization vector, effectively doubling the key space by adding an additional 128 bits of keying material.
- A signed message integrity check value that is not vulnerable to tampering or spoofing.
- An encrypted frame counter that is incorporated to thwart replay attacks.

However, because WPA uses cryptographic algorithms similar to those used by WEP, you can implement it on existing hardware with a simple firmware upgrade.

The PSK mode of WPA also allows small organizations and home office users to use a shared key WLAN without any of the vulnerabilities of static WEP. However, the viability of this option depends on choosing a pre-shared key that is strong enough to avoid simple password-guessing attacks. Like the RADIUS-based WPA and dynamic WEP, individual encryption keys are generated for each wireless client. The pre-shared key is used as an authentication credential; if you possess the key, then you are authorized to use the WLAN and receive a unique encryption key to protect the data.

The 802.11i RSN standard will bring even higher levels of security to WLANs, including better protection against denial of service attacks (DoS). The new standard was expected to be released in mid-2004.

EAP Authentication Methods

As the word “Extensible” in its name implies, EAP supports many authentication methods. These methods can use different authentication protocols such as the Kerberos version 5 authentication protocol, the Transport Layer Security (TLS) protocol, and the Microsoft–Challenge Handshake Authentication Protocol (MS–CHAP). They can also use a range of credential types that include passwords, certificates, one-time password tokens, and biometrics. Although you can theoretically use any EAP method with the 802.1X protocol, not all of the methods are suitable for use with WLANs. In particular, the method you choose must be suitable for use in an unprotected environment and be able to generate encryption keys.

The principal EAP methods in use for WLANs are EAP–TLS, Protected EAP (PEAP), Tunneled TLS (TTLS), and Lightweight EAP (LEAP). Of these, PEAP and EAP–TLS are supported by Microsoft.

EAP–TLS

EAP–TLS is an IETF standard (RFC 2716) and is probably the most widely supported authentication method on both wireless clients and RADIUS servers in use today. The EAP-TLS method uses public key certificates to authenticate both the wireless clients and the RADIUS servers by establishing an encrypted TLS session between them.

PEAP

PEAP is a two stage authentication method. The first stage establishes a TLS session to the server and allows the client to authenticate the server using the server's digital certificate. The second stage requires a second EAP method tunneled inside the PEAP session to authenticate the client to the RADIUS server. This allows PEAP to use a variety of client authentication methods including passwords with the MS–CHAP version 2 (MS–CHAP v2) protocol, and certificates using EAP–TLS tunneled inside PEAP. The EAP methods such as MS–CHAP v2 are not secure enough to be used without PEAP protection because without it they would be vulnerable to offline dictionary attacks. Support for PEAP is widespread in the industry, and Microsoft Windows XP SP1 and Pocket PC 2003 have built-in support for PEAP.

TTLS

TTLS is a two stage protocol similar to PEAP that uses a TLS session to protect a tunneled client authentication. Besides tunneling EAP methods, TTLS can also use non–EAP versions of authentication protocols such as CHAP, MS–CHAP, and others. Microsoft and Cisco do not support TTLS, although TTLS clients for a number of platforms are available from other vendors.

LEAP

LEAP is a proprietary EAP method developed by Cisco that uses passwords to authenticate clients. Although popular, LEAP only works with hardware and software from Cisco and a few other vendors. LEAP also has several published security vulnerabilities, such as susceptibility to offline dictionary attacks (which may allow attackers to discover users' passwords) and man-in-the-middle attacks. In a domain environment, LEAP can only authenticate the *user* to the WLAN, not the *computer*. Without computer authentication, machine group policies will not execute correctly, software installation settings, roaming profiles, and logon scripts may all fail, and users cannot change expired passwords.

There are WLAN security solutions that use the 802.1X protocol with other EAP methods. Some of these EAP methods, such as EAP–MD5, have significant security weaknesses when used in a WLAN environment. For this reason, they should never be used. There are other methods that support the use of one-time password tokens and other authentication protocols, such as the Kerberos protocol. However, these have yet to make a significant impact on the WLAN market.

The Benefits of 802.1X with WLAN Data Protection

In summary, the key benefits of an 802.1X protocol-based solution for your WLAN are:

- **High security:** The protocol provides a high security authentication scheme because it can use client certificates or user names and passwords.
- **Stronger encryption:** The protocol allows strong encryption of network data.

- **Transparent:** The protocol provides transparent authentication and connection to the WLAN.
- **User and computer authentication:** The protocol allows separate authentication methods for the users and computers in your environment. Separate computer authentication allows you to manage the computers in your environment even when no users are logged on to them.
- **Low cost:** Inexpensive network hardware.
- **High performance:** Because encryption is performed in WLAN hardware and not by client computer CPU, WLAN encryption has no impact on the performance level of the client computer.

There also are a few caveats to an 802.1X protocol-based solution.

- Although the 802.1X protocol has gained near universal acceptance, the use of different EAP methods means that interoperability is not always guaranteed.
- WPA is still in the early stages of adoption and may not be available on older hardware.
- The next generation RSN (802.11i) standard is not yet ratified and will require deployment of hardware and software updates (network hardware will typically need a firmware update).

However, these are relatively minor issues and are easily outweighed by the benefits of the 802.1X protocol; particularly when weighed against the serious shortcomings of the alternative approaches that are discussed later in this chapter.

The Resilience of the 802.1X Solution to Security Threats

The principal security threats to WLANs were described earlier in a table in this chapter. The following table reassesses the threats posed against a solution based on the 802.1X protocol and WLAN data protection.

Table 2.2: Security Threats Assessed Against the Proposed Solution

Threat	Mitigation
Eavesdropping (disclosure of data)	Dynamically assigning and changing encryption keys at frequent intervals and the fact that keys are unique to each user session means that as long as the key refresh is sufficiently frequent, discovering the keys and accessing data is not possible by any currently known means. WPA brings greater security by changing encryption keys per packet. Global key (protecting broadcast traffic) is rekeyed per packet.
Interception and modification of transmitted data	Enforcing data integrity and strong data encryption between the wireless client and the wireless AP ensures that it is infeasible for a malicious user to intercept and modify data in transit. Mutual authentication between the client, the RADIUS server, and the wireless AP makes it difficult for any of these to be impersonated by an attacker. WPA improves data integrity with the Michael protocol.
Spoofing	Secure authentication to the network prevents unauthorized individuals from connecting to the network and introducing spoofed data from the inside.

(continued)

DoS	Data-flooding and other DoS attacks at network level are prevented by controlling access to the WLAN using the 802.1X protocol. There is no defense against low level 802.11 DoS attacks in either dynamic WEP or WPA. This is being addressed by the 802.11i standard. However, even this new standard will not be immune to physical layer (radio-level) disruption of networks. These vulnerabilities are a feature of current 802.11 WLANs and common to all the other options discussed later in this chapter.
Free-loading (resource theft)	Unauthorized use of the network is prevented by the requirement for strong authentication.
Accidental threats	Accidental connection to the WLAN is prevented by the requirement for secure authentication.
Rogue WLANs	Although the solution does nothing directly to deal with rogue wireless APs, implementing a secure wireless solution such as this largely takes away the motivation for setting up unofficial WLANs. However, you should plan on creating and publishing a clear policy prohibiting the use of unapproved WLANs. You can enforce the policy by using software tools that scan the network for wireless AP hardware addresses, and by using handheld WLAN detection equipment.

Other Approaches to WLAN Security

The previous section discussed 802.1X authentication with WLAN data protection in some detail. This section details the other four alternatives to WLAN security listed earlier (in the beginning of the “How to (Really) Secure Your WLAN” section).

The four other approaches listed were:

- Not to deploy WLAN technology
- Stay with 802.11 static WEP security
- Use VPN to protect data on the WLAN
- Use IPsec to protect WLAN traffic

The key differentiators between these approaches and an 802.1X protocol-based solution are summarized in the following table (although the “No WLAN” option is not included because it is not directly comparable with the others). These options are covered in greater detail in later sections in this chapter.

Table 2.3: Comparison of WLAN Security Approaches

Feature	802.1X WLAN	Static WEP	VPN	IPsec
Strong authentication (1)	Yes	No	Yes, but not VPNs using shared key authentication	Yes, if using certificate or Kerberos authentication

(continued)

Strong data encryption	Yes	No	Yes	Yes
Transparent connection and reconnection to WLAN	Yes	Yes	No	Yes
User authentication	Yes	No	Yes	Yes
Computer authentication(2)	Yes	Yes	No	Yes
Broadcast and multicast traffic protected	Yes	Yes	Yes	No
Additional network devices required	RADIUS servers	No	VPN servers, RADIUS servers	No
Secures access to the WLAN itself	Yes	Yes	No	No

(1) Many VPN implementations that use IPsec tunnel mode employ a weak, shared key authentication scheme known as *XAuth*.

(2) Computer authentication means that the computer will stay connected to the WLAN and the corporate network even when no user is logged on to the computer. This capability is needed for the following Windows domain features to work properly:

- Roaming user profiles
- Computer Group Policy settings (particularly startup scripts and deployed software)
- User logon scripts and software deployed using Group Policy

Alternative 1: Not Deploying WLAN Technology

Perhaps the most obvious way of dealing with WLAN security threats is to avoid them altogether by not deploying any WLANs. Besides foregoing the benefits of WLANs outlined earlier in this chapter, this strategy is not free of pitfalls. Organization taking this approach must deal with what the META Group calls the “Price of Postponement,” which is more than just an opportunity cost. The META Group produced a study on the unmanaged way in which the use of wired LANs grew in many organizations over a decade ago. In most cases, central IT departments were forced to step in and take control of the LAN deployment reactively. Typically, the cost of re-engineering the multitude of independent and often incompatible departmental LANs was huge. For more information, see the article “How Do I Limit My Exposure Against the Wireless LAN Security Threat? The New Realities of Protecting Corporate Information,” published by the META Group on December 12, 2002.

This same threat has resurfaced with WLANs, especially in larger organizations where it is impossible to physically see what is happening in each location. Unmanaged grassroots deployment of WLANs, made possible by the extremely low cost of the components, is potentially the worst scenario. This exposes the organization to all the security threats outlined earlier, but without the central IT group knowing anything about it or being able to take steps to combat the threats.

For these reasons, if your strategy is to not adopt WLAN technology, you need to pursue this strategy in an active rather than a passive way. You should back up this decision with a clear, published policy and ensure that all employees are aware of it, as well as the consequences of violating it. You may also want to consider using scanning equipment and network packet monitors to detect the use of unauthorized wireless equipment on your network.

Alternative 2: Use 802.11 Basic Security (Static WEP)

Basic 802.11 security (static WEP) uses a shared key to control access to the network and uses the same key to encrypt wireless traffic. This simple authorization model is often supplemented with port filtering based on WLAN card hardware addresses, although this is not part of 802.11 security. The main attraction of this approach is its simplicity. Although it provides some level of security over an unsecured WLAN, this approach has serious management as well as security drawbacks, particularly for larger organizations.

The drawbacks of using static WEP include the following:

- Static WEP keys can be discovered in a matter of hours on a busy network using a computer with a WLAN adapter and hacking tools such as Aircrack-ng or WEPCrack.
- The most serious weakness of static WEP is that there is no mechanism for dynamically assigning or updating the network encryption key. Without 802.1X and EAP to enforce regular key updates, the encryption algorithm that static WEP uses is vulnerable to key recovery attacks.
- The static keys can be changed, but the process for doing this on the APs and wireless clients is usually manual and always time consuming. Moreover, the keys must be simultaneously updated on the clients and the APs to preserve the clients' connectivity. In practice, this is so difficult to achieve that the keys are usually left unchanged.
- The static key needs to be shared between all users of the WLAN and all wireless APs. This situation creates a vulnerability because a secret shared between a large number of people and devices is unlikely to remain a secret for long.

Static WEP provides WLANs with a very limited access control mechanism based on knowing the WEP key. If you discover the name of the network, which is easy to do, as well as the WEP key, you can connect to the network.

One way of improving this situation is to configure the wireless APs to allow only a predefined set of client network adapter addresses. This is commonly known as media access control (MAC) address filtering; the MAC layer refers to the low-level firmware of the network adapter.

Network adapter address filtering for access control has its own set of issues:

- Manageability is extremely poor. Maintaining a list of hardware addresses for anything but a small number of clients is difficult. Also, distributing and synchronizing this list across all your APs is a significant challenge.
- Scalability is poor. The APs have a finite filter table size limit, thus restricting the number of clients that you can support.
- There is no way to associate a MAC address with a user name, so you can only authenticate by computer identity and not user identity.
- An intruder could spoof an "allowed" MAC address. If a legitimate MAC address is discovered, it is easy for an intruder to use this address instead of the predefined address burned onto the adapter.

Pre-shared key solutions are only practical for small numbers of users and APs due to the difficulty of managing key updates across multiple locations. Cryptographic flaws with WEP make its usefulness extremely questionable, even in very small environments.

Conversely, WPA's pre-shared key mode does provide a good level of security with very low infrastructure overhead for small organizations. A wide range of hardware also supports WPA PSK, and you can manually configure WLAN clients. For these reasons, WPA PSK is the configuration of choice for SOHO environments.

Alternative 3: Use Virtual Private Networks

VPNs are probably the most popular form of network encryption; a lot of people rely on the tried and trusted VPN technologies to protect the confidentiality of data sent over the Internet. When the vulnerabilities of static WEP were discovered, VPN was quickly proposed as the best way to secure data traveling over a WLAN. This approach was endorsed by analysts such as the Gartner Group and, unsurprisingly, VPN solution vendors enthusiastically promoted it.

VPN is an excellent solution to securely traverse a hostile network such as the Internet (although the quality of VPN implementations varies). However, it is not necessarily the best solution for securing internal WLANs. For this kind of environment, a VPN offers little or no additional security compared with 802.1X solutions, but it does significantly increase complexity and costs, reduce usability, and render important pieces of functionality inoperable.

Note: These limitations are distinct from using a VPN to secure traffic over public wireless LAN hotspots. Protecting the network data of users connecting over hostile remote networks is a legitimate use of a VPN. In this kind of scenario users expect secure connectivity to be more intrusive and less functional than a LAN connection; something that they do not expect when inside the company's own premises.

The advantages of using VPNs to protect WLANs include the following:

- Most organizations already have a VPN solution deployed so users and IT staff will be familiar with the solution.
- VPN data protection normally uses software encryption that allows algorithms to be changed and upgraded more easily than hardware-based encryption.
- You may be able to use relatively less expensive hardware because VPN protection is independent of WLAN hardware (although the price premium for 802.1X capable network hardware is nearly gone).

The disadvantages of using VPNs in place of native WLAN security include:

- VPN lacks user transparency. VPN clients usually require the user to manually initiate a connection to the VPN server. Therefore, the connection will never be as transparent as a wired LAN connection. For Non-Microsoft VPN clients, in addition to the standard network or domain logon, they may also be prompted for logon credentials when they attempt to connect to the network. If the VPN disconnects them because of a poor WLAN signal or because the user is roaming between APs, the clients must reconnect to the network.
- Because only the user initiates the VPN connection, an idle, logged-off computer will not be connected to the VPN (and thus the corporate LAN). Therefore, a computer cannot be remotely managed or monitored unless a user is logged on to it. This can prevent certain computer Group Policy object (GPO) settings, such as startup scripts and computer assigned software, from being applied.

- Roaming profiles, logon scripts, and software deployed to the user using GPOs may not work as expected. Unless the user chooses to log on using the VPN connection from the Windows logon prompt, the computer will not connect to the corporate LAN until after the user has logged on and initiated the VPN connection. Attempts to access the secure network before this will fail. With a non-Microsoft VPN client, it may be impossible to do a full domain logon over the VPN connection.
- Resuming from standby or hibernation does not automatically re-establish the VPN connection; the user has to do this manually.
- Although the data inside the VPN tunnel is protected, the VPN offers no protection for the WLAN itself. An intruder can still connect to the WLAN and attempt to probe or attack any devices attached to it.
- The VPN server(s) can become a constraint. All WLAN client access to the corporate LAN is channeled through the VPN server. VPN devices traditionally service a large number of relatively low speed remote clients. Therefore, most VPN gateways cannot handle tens or hundreds of clients running at full LAN speed.
- The cost of additional hardware and ongoing management of the VPN devices is likely to be much higher than a native WLAN solution. Each site will typically need its own VPN server in addition to WLAN APs.
- VPN sessions are more prone to being disconnected when clients roam between APs. Although applications will often tolerate a momentary disconnection when switching wireless APs, even a brief interruption of a VPN session will often require the user to manually reconnect to the network.
- The cost of VPN server and client software licenses, as well as the cost of deploying the software, may be an issue with non-Microsoft VPN solutions. You may also have concerns with the VPN client software compatibility because non-Microsoft clients often replace core Windows functionality.
- Many analysts and vendors make the assumption that VPN security is always better than WLAN security. Though this may be true for static WEP, it is not necessarily the case for the 802.1X EAP-based solutions described in this chapter. In particular, VPN authentication methods are often far less secure and are at best unlikely to be significantly stronger. For example, the WLAN solutions supported by Microsoft use exactly the same EAP authentication methods as its VPN solutions (EAP-TLS and MS-CHAP v2). Many VPN implementations, especially those based on IPsec tunnel mode, use pre-shared key authentication (a group password). This has been widely discredited and shown to have serious security vulnerabilities, ironically, sharing some of these vulnerabilities with static WEP.
- A VPN does nothing to secure the WLAN itself. Though the data inside the VPN tunnels is secure, anyone can still connect to the WLAN and attempt to attack legitimate wireless clients and other devices on the WLAN.

VPN is ideally suited to securing traffic passing over hostile networks, whether the user is connecting over a home broadband connection or from a wireless hotspot. However, VPN was never designed to secure network traffic on internal networks. Because of this, for most organizations, VPN in this role is too cumbersome, and the functionality is too limiting for the user and too costly and complex for the IT department to maintain.

In exceptional cases where higher security for a particular connection or traffic type is needed, this can be provided by a VPN tunnel or IPsec transport mode *in addition* to the native WLAN protection. This is a more sensible use of network resources.

Alternative 4: Use IP Security

IPsec allows two network peers to securely authenticate each other and authenticate or encrypt individual network packets. You can use IPsec to either securely tunnel one network over another or simply to protect IP packets being transmitted between two computers.

IPsec tunneling is typically used in client access or site-to-site VPN connections. IPsec tunnel mode is a form of VPN that works by encapsulating a whole IP packet within a protected IPsec packet. Like other VPN solutions, this adds an overhead to the communication that is not really required for communication between systems on the same network. The pros and cons of IPsec tunnel mode were covered in the discussion on VPN in the previous section.

IPsec can also secure end-to-end traffic between two computers (without tunneling) using IPsec *transport mode*. Like VPN, IPsec is an excellent solution in many circumstances, although it is not a replacement for native WLAN protection that you can implement at the hardware layer.

Some of the advantages of IPsec transport mode protection are:

- It is transparent to users. Unlike VPNs, no special logon procedure is required.
- IPsec protection is independent of WLAN hardware. It only requires an open, unauthenticated WLAN. Unlike VPN, no additional servers or devices are required because the security is negotiated directly between the computers at each end of the communication.
- Use of cryptographic algorithms is not constrained by the WLAN hardware.

The disadvantages of using IPsec in place of native WLAN security are:

- IPsec uses computer-level authentication only; there is no way to implement a user-based authentication scheme along with it. For many organizations, this will not be a problem but it does allow *unauthorized* users to connect to other IPsec protected computers on the network if they manage to log on to an *authorized* computer.

Note: Some IPsec implementations on non–Windows platforms rely on user-only authentication. However, as with the VPN solution, the computer will not be connected to the network when the user is not logged in, thus preventing certain management operations and disabling user setting functionality.

- Managing IPsec policies can be complex for a large organization. Attempts to enforce general IP traffic protection may interfere with more specialized uses of IPsec where end-to-end protection is required.
- Full security requires encrypting all end-to-end traffic, but some devices may not be capable of using IPsec. This will force traffic to these devices to be transmitted unencrypted. IPsec will provide no protection to these devices, exposing them to anyone who connects to the WLAN.
- Because IPsec protection occurs at the network level rather than at the MAC layer, it is not fully transparent to network devices such as firewalls. Some IPsec implementations will not work across a network address translation (NAT) device.
- End-to-end IPsec cannot protect broadcast or multicast traffic because IPsec relies on two parties mutually authenticating and exchanging keys.

- Although the data inside the IPsec packets is protected, the WLAN itself is not protected. An intruder can still connect to the WLAN and attempt to probe or attack any devices connected to it or listen to any traffic that is not protected by IPsec.
- IPsec network traffic encryption and decryption increases the load on computer CPUs. This can overload heavily used servers. Although this processing overhead can be offloaded to specialized network cards, most servers are not normally equipped with these cards.

Like VPN, IPsec is an excellent solution for many security scenarios, but it does not address WLAN security as well as native WLAN protection.

Selecting the Right WLAN Options

Based on the discussion in the previous section, the 802.1X WLAN solution is by far the best of the available alternatives. However, as noted in the “Understanding WLAN Security” section, once you have decided to use an 802.1X solution, you have to choose from a number of options to make the solution work.

The two key choices are:

- Whether to use passwords or certificates to authenticate your users and computers.
- Whether to use dynamic WEP or WPA WLAN data protection.

These two choices are independent of each other.

As discussed earlier in this chapter, Microsoft has two WLAN security solution guides; one focused on using password authentication and the other on using certificate authentication. Both solutions work with either dynamic WEP or WPA.

Deciding On the Right WLAN Security Solution

The following flowchart summarizes the choices between the two WLAN security solution guides.

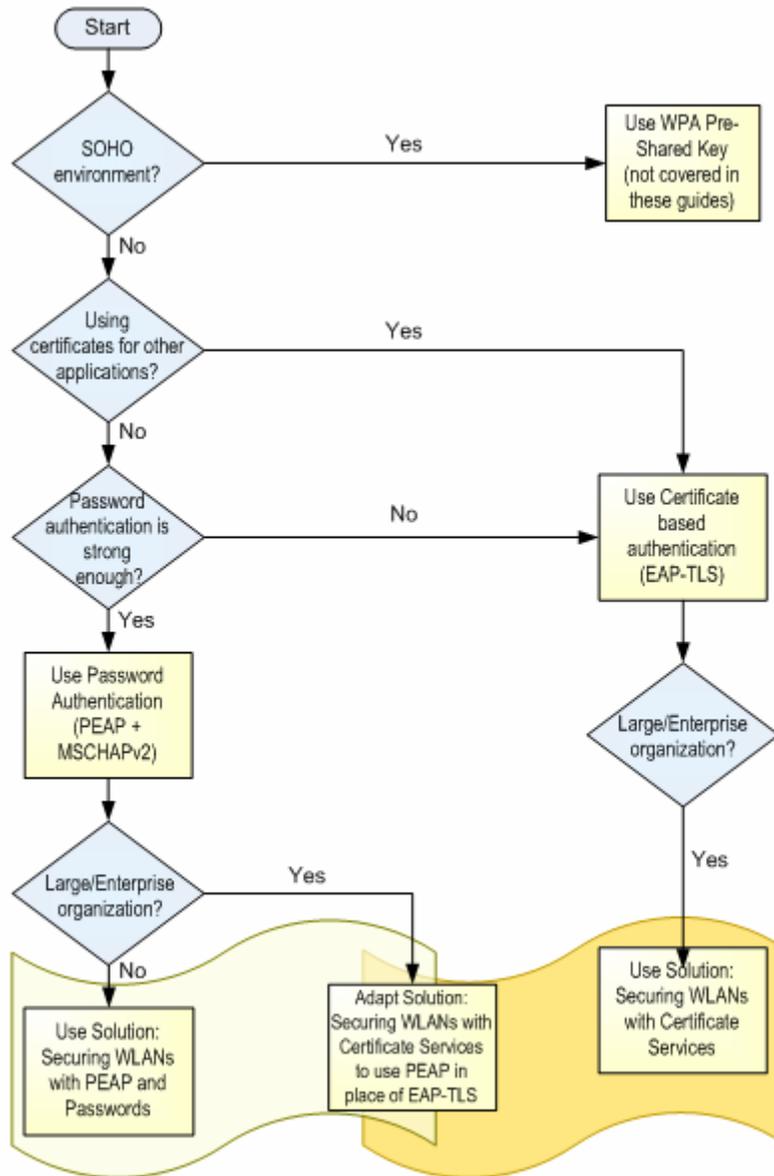


Figure 2.2

The decision tree for the WLAN security solutions

The outcome of this decision tree depends on the size and specific security requirements of your organization. Most organizations can use one or the other of the Microsoft WLAN solutions without any modification. For example, most small-to-medium organizations will choose the simpler password-based authentication solution described in the *Securing WLANs with PEAP and Passwords* solution guide. Larger organizations are more likely to move toward using the digital certificate-based *Securing Wireless LANs with Certificate Services* solution guide.

Although each solution was written with these audiences in mind, there is a good deal of latitude with each solution. *Securing Wireless LANs with PEAP and Passwords* can be deployed in organizations ranging in size from just tens of users to many thousands of users. The *Securing Wireless LANs Certificate Services* solution applies to organizations ranging in size from a few hundred to tens of thousands of users (organizations with

fewer than five hundred users normally do not have sufficient IT resources to deploy and maintain certification authorities).

One common case that isn't directly covered by either of the guides is large organizations deploying a password-based WLAN solution. Although the technical detail in the *Securing Wireless LANs with PEAP and Passwords* solution is equally applicable to large and small businesses, much of the design, planning, and operational detail required by larger organizations has been omitted in the interest of simplicity. Fortunately, the similarity between the architecture and technical components used in both solutions allows you to mix-and-match parts of the solutions relatively easily. The *Securing Wireless LANs with PEAP and Passwords* solution has an appendix that gives you some guidance on which parts from each solution are relevant to large organizations that wish to deploy a password-based WLAN solution.

Choosing Between Dynamic WEP and WPA

WEP data protection, when combined with the strong authentication and dynamic key update that 802.1X and EAP offer, provides you with a level of security that is more than adequate for most organizations. However, the WPA standard improves on this and provides even better levels of security.

The differences between using WPA and a dynamic WEP in either of the solutions are minimal, and migrating from a dynamic WEP environment to a WPA environment is very simple. The key changes involved in moving from dynamic WEP to WPA are:

- You must obtain and deploy firmware updates for your network hardware (wireless APs and wireless network adapters) if the hardware does not currently support WPAit. Firmware updates for wireless network adapters are often included in network driver updates.
- You must enable WPA on your wireless APs.
- You must change the WLAN client configuration to negotiate WPA instead of WEP security.
- You should increase the session time-out on the Internet Authentication Service (IAS) remote access policy, which is used to force WEP key refresh, to reduce the load on the IAS server.

Note: IAS is the Microsoft RADIUS server implementation. It is included in Windows Server 2003 but not installed by default.

WPA should be your first choice, if it is available to you. However, consider whether any of the following issues will make using WPA more problematic:

- Your network hardware may not yet support WPA (this is unlikely with new devices, but you may have a large installed base of pre-WPA hardware).
- Support for GPO controlled settings is available only in the next update to Windows Server 2003; other versions do not have this support and you must manually configure WPA settings on Windows XP clients.
- WPA may not be supported on all of your clients; for example, Windows 2000 and earlier and Pocket PC currently have no built-in support for WPA.

If you decide that you are not yet in a position to deploy WPA, you should deploy a dynamic WEP solution and plan to migrate to WPA when circumstances permit.

Summary

This chapter provides information you can use to define a wireless LAN security strategy for your organization. The first part of the chapter explored the business advantages of wireless networks, and the security threats confronted by poorly protected WLANs. The middle of the chapter examined how wireless LAN security based on the 802.1X protocol, EAP, and strong data protection works to mitigate these threats. The relative merits of alternatives such as VPNs, IPsec, and static WEP security were also discussed. The last part of the chapter included guidance on how to determine which of the WLAN security options is best for your organization, and which of the two Microsoft WLAN security solutions is most appropriate for your organization.

References

This section provides references to important supplementary information and other background material relevant to this chapter.

- The Microsoft solution for [Securing Wireless LANs with PEAP and Passwords](http://go.microsoft.com/fwlink/?LinkId=23459) is available at <http://go.microsoft.com/fwlink/?LinkId=23459>.
- For more detailed technical information about IEEE 802.11 and related technologies, see the “[802.11 Wireless Technical Reference](#)” section of the Windows Server 2003 Technical Reference available at www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/techref/w2k3tr_wir_intro.asp.
- For more information about 802.11, see the [IEEE 802.11](http://www.ieee802.org/11/) page of the IEEE 802.11 Standards News Bulletin at www.ieee802.org/11/.
- For more information about 802.1X, see the [802.1x - Port Based Network Access Control](http://www.ieee802.org/1/pages/802.1x.html) page at www.ieee802.org/1/pages/802.1x.html.
- For more information about the EAP standard, see [RFC 2284](http://www.ietf.org/rfc/rfc2284.txt?number=2284) at www.ietf.org/rfc/rfc2284.txt?number=2284.
- For more information about the Wi-Fi Alliance WPA standard, see the [Wi-Fi Alliance Overview](http://www.wi-fi-allyance.org/OpenSection/pdf/Wi-Fi_Protected_Access_Overview.pdf) at www.wi-fi-allyance.org/OpenSection/pdf/Wi-Fi_Protected_Access_Overview.pdf.
- For more information about wireless networking, see the [Wi-Fi](http://www.microsoft.com/wifi) page of the Microsoft Windows Server System Web site at www.microsoft.com/wifi.
- For a detailed discussion of PEAP and how it compares with LEAP (and also EAP-TLS and EAP-MD5), see “[The Advantages of Protected Extensible Authentication Protocol \(PEAP\): A Standard Approach to User Authentication for IEEE 802.11 Wireless Network](#),” article at www.microsoft.com/windowsserver2003/techinfo/overview/peap.mspx.
- The META Group article “[How Do I Limit My Exposure Against the Wireless LAN Security Threat? The New Realities of Protecting Corporate Information](#)” at www.metagroup.com/cgi-bin/inetcgi/jsp/displayArticle.do?oid=35725.

3

Secure Wireless LAN Solution Architecture

Introduction

The previous chapter discussed the options for wireless local area network (WLAN) security and described why 802.1X wireless authentication using the Extensible Authentication Protocol–Transport Layer Security (EAP-TLS) protocol was selected for this solution. This chapter describes the solution architecture, and then a logical design is derived based on design criteria taken from an example company. You can use this information as the foundation to implement the solution. The logical design is based on 802.1X WLAN network hardware, Remote Authentication Dial-In User Service (RADIUS) authentication, and a Public Key Infrastructure (PKI).

Chapter Prerequisites

You should have an understanding of IT infrastructure design concepts, as well as a familiarity with the key components that form part of the design. The key components are: WLANs and networking components, RADIUS, the Active Directory® directory service, and PKIs. Detailed knowledge of these items is not required.

Chapter Overview

This chapter aims to:

- Provide a conceptual overview of how a secure WLAN solution based on the 802.1X and EAP–TLS protocols functions and the key components of this type of solution.
- Define the solution design criteria for the logical design and the later stages of the detailed technical design.
- Produce a coherent logical design that will form the basis for the detailed design in the subsequent chapters.
- Explain how you can scale the solution to meet the needs of organizations of different sizes.
- Detail some of the ways in which you can extend the proposed design or use it as the basis to build other network access solutions that include virtual private networks (VPN) and wired network access control, and examine how you can use the PKI component of the design as the foundation for a variety of security applications.

Subsequent chapters will look at the detailed design process for each of the main components of the logical design (WLAN, RADIUS, and PKI) in preparation to build and operate the solution.

Conceptual Design

As discussed in the previous chapter, there are a number of serious security vulnerabilities inherent in wireless networking. At best, these weaknesses are only partially addressed by the use of Wired Equivalent Privacy (WEP) as specified in the Institute of Electrical and Electronics Engineers (IEEE) 802.11 standard. The solution proposed in this guide addresses the problem of how to improve the security of wireless network communications. To do this, the ideal solution needs to have the following features:

- Robust wireless client authentication. This should include mutual authentication of the client, the wireless access point (AP), and the RADIUS server.
- An authorization process to determine who will be allowed to access the wireless network.
- Access control to only permit network access to authorized clients.
- Strong encryption of wireless network traffic.
- Secure management of encryption keys.
- Resilience to denial of service (DoS) attacks.

The 802.1X protocol standard for network access control combined with a secure authentication method such as EAP-TLS fulfills some of these requirements. High strength WEP provides relatively secure encryption of network traffic, but poor key management. Methods to manage WEP encryption keys inherent in 802.1X and EAP are much more secure than the 802.11 base standards permit. The WiFi Protected Access (WPA) standard is a collection of industry-based standards that includes 802.1X and EAP (among other improvements), and a standardized protocol for key management called Temporal Key Integrity Protocol (TKIP). The WPA standard represents a considerable step forward for WLAN security and has been endorsed by most analysts and vendors.

Note: None of the WPA improvements address some of the DoS weaknesses inherent in both 802.11 and 802.1X. The DoS weaknesses are not as serious as any of the other WEP flaws, and almost all of the demonstrated DoS attacks cause only temporary network disruption. However, the threat of DoS attacks are still a serious concern for some organizations, and one that is unlikely to be resolved before the release of the IEEE 802.11i standard that is anticipated in 2004.

Although support for WPA is now widespread, there are still many existing devices and systems that do not support it. For this reason, the solution for this guide is designed to work with both dynamic WEP and WPA. Most network hardware vendors sell products that support 802.1X with dynamic WEP keys and WPA. For the purposes of this design chapter, the two approaches are treated interchangeably — the use of one instead of the other has no significant impact on the design.

A conceptual figure of the solution (802.1X EAP-TLS Authentication) is displayed in the following figure.

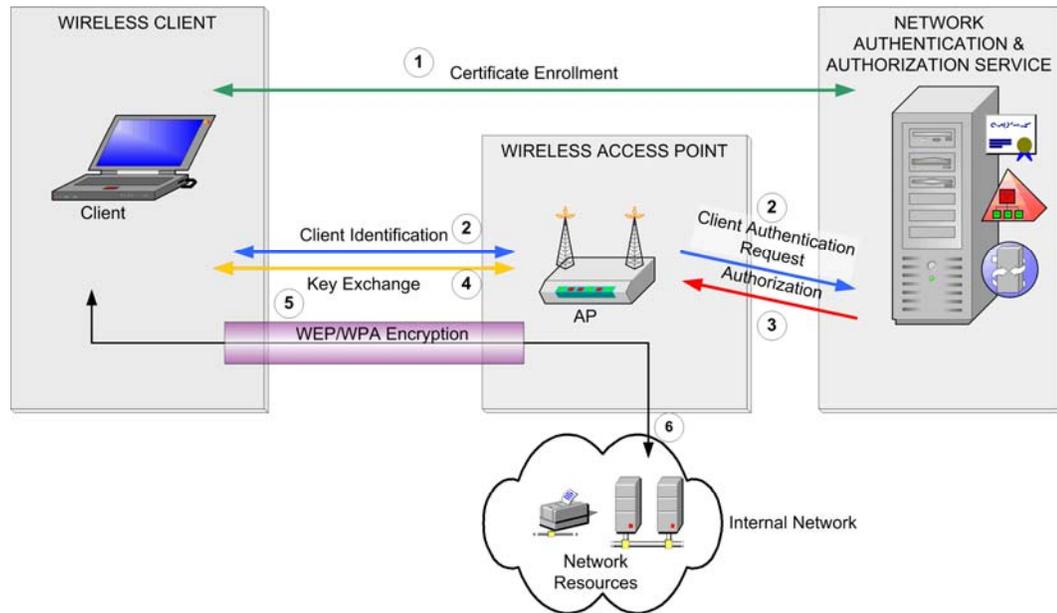


Figure 3.1

The solution concept based on 802.1X EAP-TLS authentication

The figure depicts four main components:

- The wireless client.** This component is a computer or device running an application that requires access to network resources. The client has the capability of encrypting its network traffic and of storing and securely exchanging credentials (such as keys or passwords).
- The wireless AP.** In general networking terms, this component is known as the Network Access Service (NAS), but wireless standards refer to this as the AP. The wireless AP implements access control functions to allow or deny access to the network and provides the capability to encrypt wireless traffic. The AP also has the means to securely share encryption keys with the client to secure network traffic. Finally, it can query an authentication and authorization service for authorization decisions.
- The Authentication Service (AS).** This component stores and verifies the credentials of valid users and makes authorization decisions based on an access policy. It may also collect accounting and audit information about client access to the network. The RADIUS server is the main component of the AS but the directory and the CA also contribute to this function.
- The internal network.** This component is a secure area of networked services that the wireless client application needs to gain access.

The numbers on the figure illustrate the network access process, which the following steps describe in more detail:

1. The wireless client must establish its credentials with the AS before wireless network access is established. (This may be accomplished using some out-of-band means—for example, by floppy disk exchange—or it may take place on a wired or other secure network.)
2. When the client computer is in range of the wireless AP, it tries to connect to the WLAN that is active on the AP. The WLAN is identified by its Service Set Identifier (SSID). The client detects the WLAN SSID and uses it to determine the correct settings and credential type to use for this WLAN.

The wireless AP is configured to allow only secured (802.1X authenticated) connections. When the client tries to connect to it, the AP issues a challenge to the client. The AP then sets up a restricted channel that allows the client to communicate only with the RADIUS server. This channel blocks access to the rest of network. The RADIUS server will only accept a connection from a trusted wireless AP, or one that has been configured as a RADIUS client on the Microsoft Internet Authentication Service (IAS) server and that provides the shared secret for that RADIUS client.

The client attempts to authenticate to the RADIUS server over the restricted channel using 802.1X. As part of the EAP–TLS negotiation, the client establishes a Transport Layer Security (TLS) session with the RADIUS server. Using a TLS session serves the following purposes:

- It allows the client to authenticate the RADIUS server; this means that the client will only establish the session with a server holding a certificate that is trusted by the client.
- It allows the client to supply its certificate credentials to the RADIUS server.
- It protects the authentication exchange against packet snooping.
- The negotiation of the TLS session generates a key that the client and RADIUS server can use to establish common master keys. These keys are used to derive the keys used to encrypt the WLAN traffic.

During this exchange, the traffic within the TLS tunnel is only visible to the client and RADIUS server and is never exposed to the wireless AP.

3. The RADIUS server validates the client credentials against the directory. If the client is successfully authenticated, the RADIUS server assembles information that allows it to decide whether to authorize the client to use the WLAN. It uses information from the directory (such as group membership) and constraints defined in its access policy (for example, the time periods in which WLAN access is allowed) to either grant or deny access to the client. The RADIUS then relays the access decision to the AP.
4. If the client is granted access, the RADIUS server transmits the client master key to the wireless AP. The client and AP now share common key information that they can use to encrypt and decrypt the WLAN traffic passing between them.
When using dynamic WEP to encrypt the traffic, the master keys need to be changed periodically to thwart WEP key recovery attacks. The RADIUS server does this by regularly forcing the client to re-authenticate and generate a new key set.

If WPA is used to secure the communication, the master key information is used to derive the data encryption keys, which are changed for each packet transmitted. WPA does not need to force frequent re-authentication to ensure key security.

5. The AP then establishes the client WLAN connection to the internal LAN, allowing the client unrestricted access to the systems on the internal network. Traffic sent between the client and AP is now encrypted.
6. If the client requires an IP address, it can now request a Dynamic Host Configuration Protocol (DHCP) lease from a server on the LAN. Once the IP address has been assigned, the client can start exchanging information normally with the systems on the rest of the network.

The following figure displays this process in more detail.

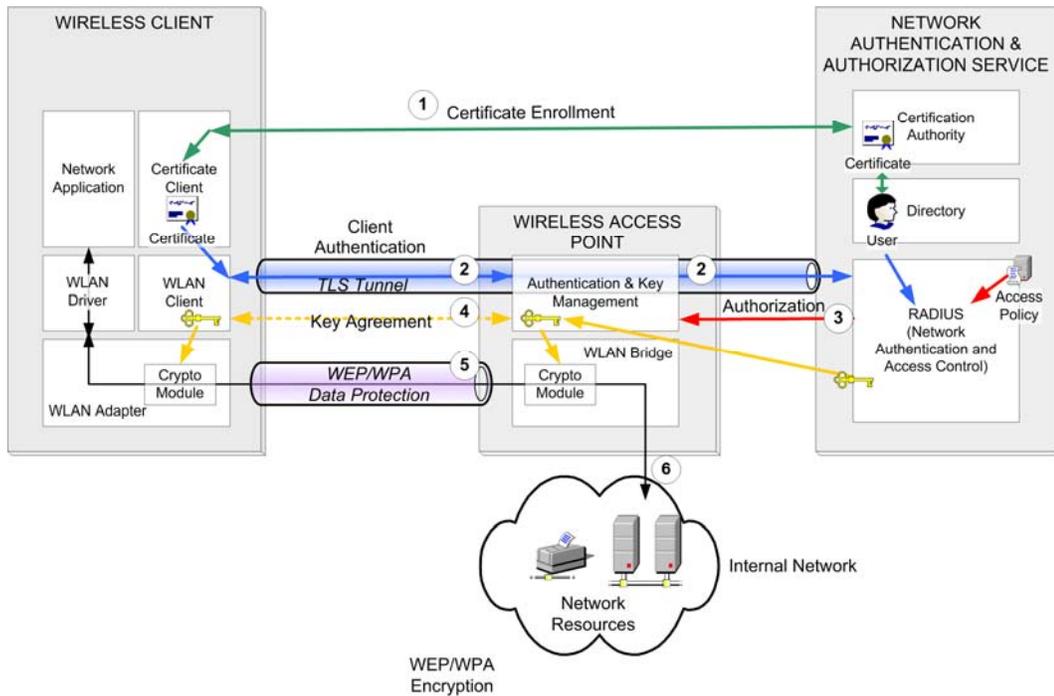


Figure 3.2

The 802.1X EAP-TLS access process

The figure displays the individual components in more detail. Later sections of this chapter will return to this diagram to expand on it. For now, you should note the subcomponents of the AS: the certification authority (CA), the directory, and the RADIUS server. Although conceptually these subcomponents perform a relatively simple set of tasks, carrying out these operations securely and in a way that is scalable, manageable, and reliable requires quite a sophisticated infrastructure. The majority of the planning, implementation, and management effort required for this is detailed in the remaining chapters of this guide.

Solution Design Criteria

Now that the the basic concepts of the solution have been described, the key design criteria for the solution can now be addressed. These criteria provide the guidelines that will turn the solution concept into a design that you can actually implement.

The design criteria are derived from the requirements of a typical organization implementing this solution. The following sections describe that organization and its main technical requirements.

Target Organization

The description of the organization in this section is intended only to provide a context for the design criteria. When assessing the suitability of the solution for your organization, focus on whether the design criteria make sense for you, not on whether your organization is exactly like the one described in this chapter.

The target organization for the solution may have deployed a WLAN in some locations to minimize network infrastructure costs, and increase staff mobility and productivity. The organization has a clear sense of the need for security and has already deployed a number of technologies to enhance its IT security. For example, it has deployed domain authentication, Internet firewalls, virus scanners, and a remote access or VPN solution. It also has longer term plans to use a number of other high security applications, such as file encryption and secure e-mail.

The simplified logical and physical network layout of this organization might look like the one defined in the following figure:

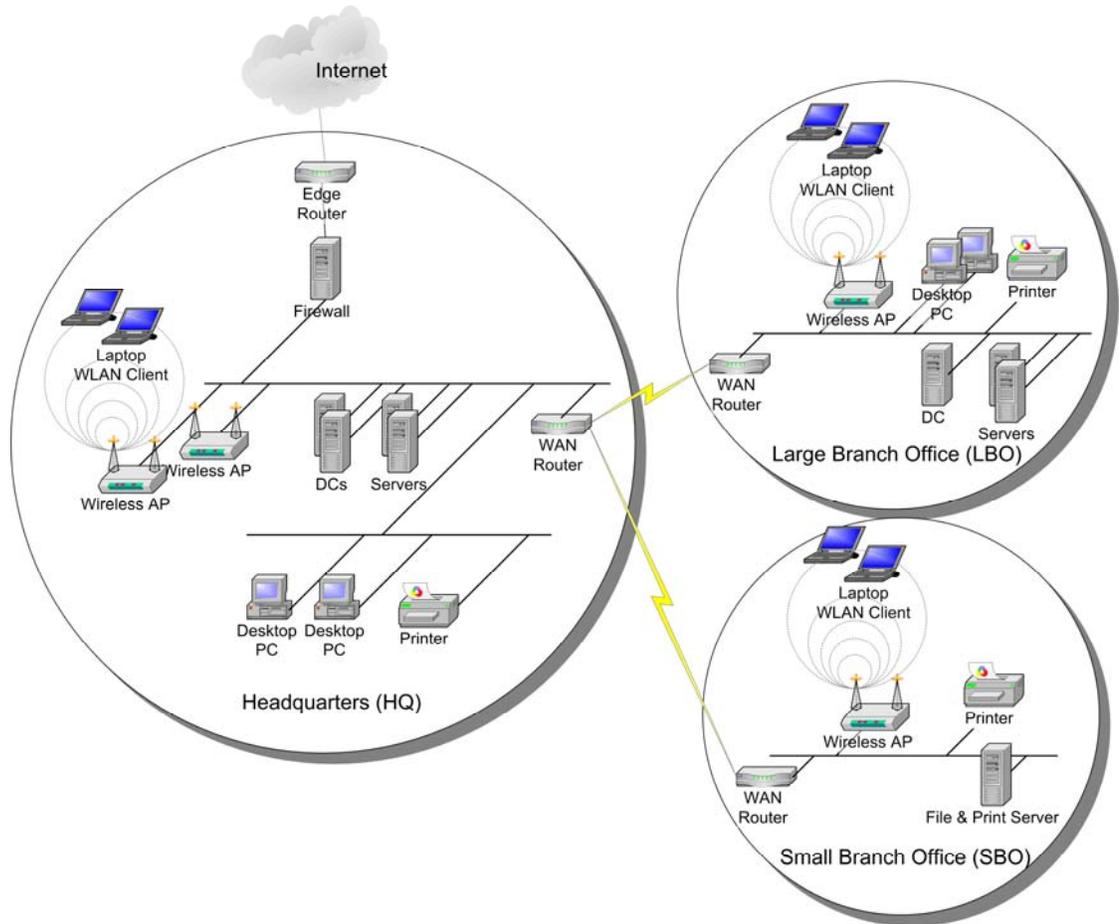


Figure 3.3

The schematic of the target organization's network and physical layout

Although only one large and one small outlying office are in the figure, in reality there could be several of each. For clarity, only a small number of servers and clients are depicted. This small number of hosts is not representative of a typical organization.

Within limits, the size of the target organization has relatively little effect on the design criteria for the solution. On the smaller end of the scale, the head office may employ a few hundred staff who work with branch offices that serve tens of staff. On the larger end of the scale, the head office staff may number in the thousands, and outlying offices may have hundreds of staff. At both ends of the scale, organizations will typically also have smaller offices with small numbers of staff.

Organization Requirements

The following requirements are typical of the organization depicted in this scenario:

- The organization must improve the security of the WLAN to eliminate or substantially reduce the following threats:
 - Intruders eavesdropping on data transmissions across the WLAN.
 - Intruders intercepting and modifying data transmissions on the WLAN.
 - Intruders or other unauthorized users connecting to the WLAN and introducing viruses or other hostile code onto the internal network.
 - Network-level (rather than radio-level) DoS attacks.
 - Intruders making use of the corporate WLAN to gain Internet access.
- The security measures should not reduce the usability of the network and not lead to any significant increase in help desk calls.
- The deployment and ongoing management costs should be low enough to justify them even if only a relatively small number of users (less than 10 percent of the workforce) use the WLAN solution.
- The design should be capable of supporting a broad variety of clients and devices.

In addition, there are usually a number of other, more general technical requirements:

- Resilience to single component failure should be maintained.
- Scalability should exist to cope with higher levels of use in the future possibly beyond 100 percent of the existing workforce . The cost to support increased numbers of users should be minimal or at least in proportion to the expansion required.
- Reusability of components—wherever possible. The solution should reuse existing infrastructure and any new components introduced by the solution should be reusable by future projects.
- Existing management and monitoring infrastructure should easily accommodate the new solution.
- The ability to recover from a catastrophic failure (for example, by restoring backups to alternative hardware) should be maintained.
- Reliance on industry standard protocols and formats should be followed. Where no current standards yet exist, the solution should align with future standards.
- Robust security (including regular renewal) of credentials and keys should be provided in the solution.
- Full audit information for user enrollment and client access to the network should be provided.

Solution Design Criteria

From these requirements, the criteria in the following table can be derived to support the solution design.

Table 3.1: Solution Design Criteria

Design factor	Criteria
Security	<ul style="list-style-type: none"> –Robust authentication and authorization of wireless clients. –Robust access control to limit network access to authorized clients. –High strength encryption of wireless network traffic. –Secure management of encryption keys. –Resilience to DoS attacks.
Scalability	Basic design that scales up and down to include a broad range of organizations.
–Min/Max users supported	<ul style="list-style-type: none"> –500–15,000or more WLAN users. –500–15,000or more certificate users.
–Number of sites supported	<ul style="list-style-type: none"> –Multiple large sites—with local authentication domain controllers and Microsoft Internet Authentication Service (IAS)—supported with resilience to wide area network (WAN) failure. –Multiple small sites supported with no resilience to WAN failure.
Component reuse (use of existing infrastructure)	Use Active Directory, network services, and Microsoft Windows® XP clients.
Component reuse (usability by future applications)	<ul style="list-style-type: none"> –Support for other network access applications (VPN and 802.1X wired network access) by the authentication infrastructure. –Support for a wide variety of applications—for example, Encrypting File System (EFS) and VPN—by PKI.
Availability	Resilience to single component or network link failure.
Extensibility	<ul style="list-style-type: none"> –Extensible to support future capabilities and standards (for example, 802.11i, WPA, 802.11a for WLAN). –Certificate infrastructure is extensible to support most common uses of public key certificates (secure e-mail, smart card logon, code signing, and Web Service Security).
Manageability	Integration into existing corporate management solutions (includes system and service monitoring, backup, configuration management).
IT organization structure	Favors centralized IT (department of at least five and typically 20–30 IT staff).
Standards conformance	Adherence to current relevant standards and a clear migration path to future relevant standards.

Solution Logical Design

This section describes the logical and logical-physical solution design. This involves specification and placement of actual components, although it does not include physical design details such as server hardware specification.

Conceptual Design Review

Using the following figure, which was shown earlier in the chapter, this section examines how the different components fit together in the overall design.

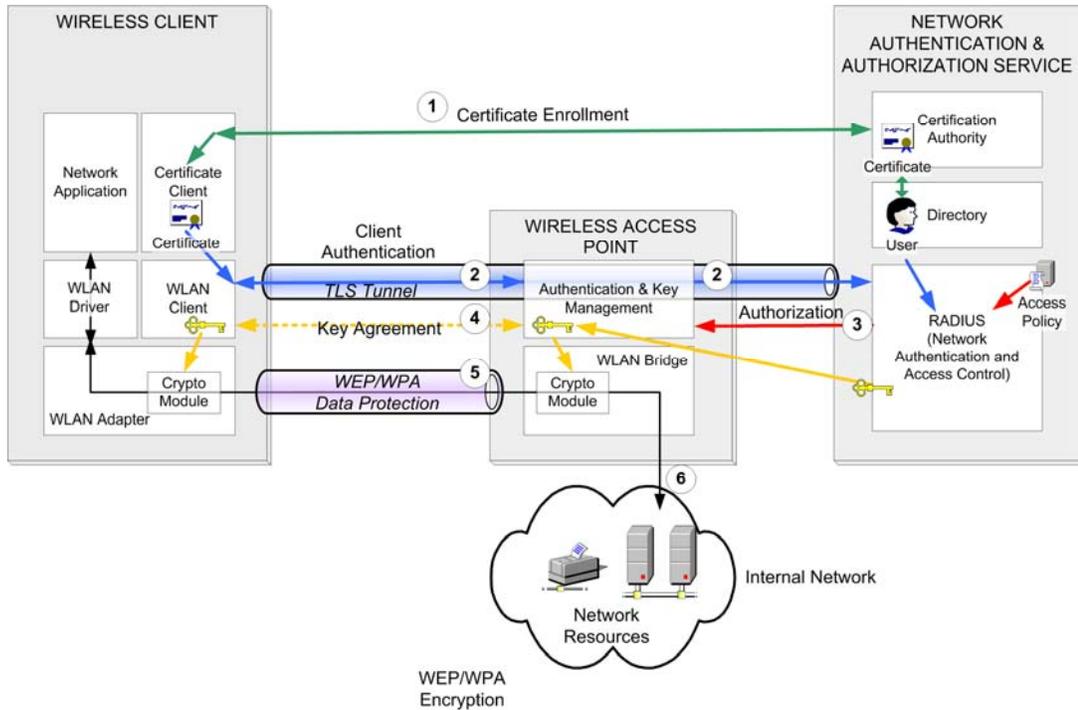


Figure 3.4
The conceptual view of the network access process

Logical Design

The previous figure divided the logical components to make it easier to understand the WLAN access process. However, to simplify deployment and management, it makes sense to slightly re-group the components.

The component grouping allows the whole design to be viewed in a modular way that allows for maximum reuse of these components. For example, it would be possible to implement the PKI component only to authenticate WLAN users. However, this may limit the reusability of the PKI component for other applications. Similarly, the RADIUS component for the solution should be designed while keeping in mind what other applications this component might be required to support in the future.

The IT services in the design are grouped into the following categories:

- WLAN components—Wireless clients and Access Points (AP)
- RADIUS component
- PKI component—Certification Authorities (CA)
- Infrastructure services component

This last component includes a directory and supporting network services. These are made up of IT services that will typically already exist in the organization that the solution interacts with in some way.

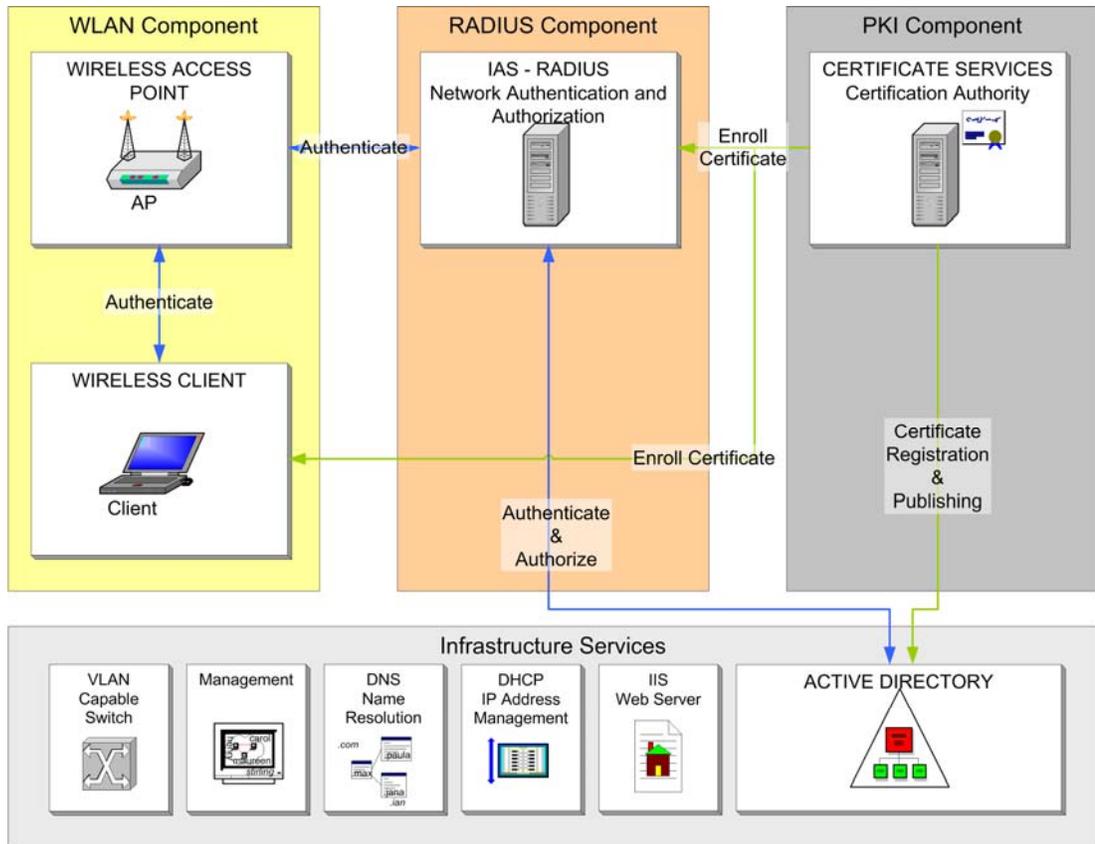


Figure 3.5
The logical design of the secure WLAN solution

Logical-Physical Level

At the logical-physical level, the design now shows how these components will be implemented as physical servers, how they will be linked together, and how they will be distributed between the different sites of the target organization. However, the numbers of servers displayed in the following figure is a generalization. The final definition of the number and placement of the servers will be discussed in the later planning chapters of this guide.

Headquarters

The following figure illustrates the server implementation in the head office. Only the top three components represent new servers or components that must be acquired. The infrastructure services components already exist in some form for many organizations. If your organization has already deployed 802.1X-capable WLAN equipment, the WLAN component may also already exist.

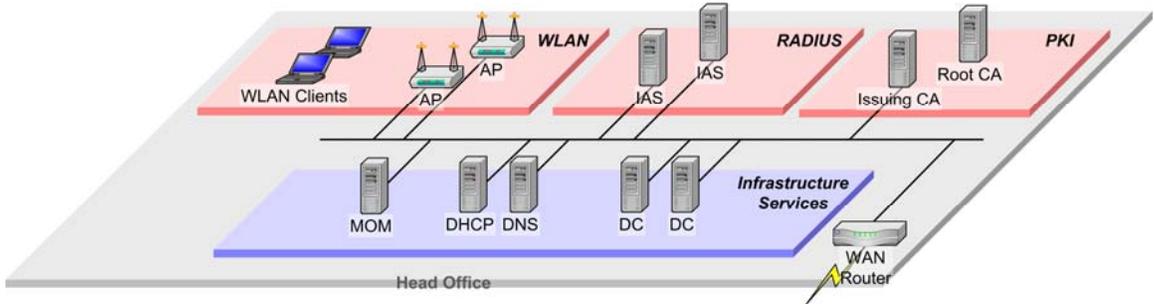


Figure 3.6
The head office server implementation

Large Branch/Regional Office

The following figure illustrates the physical layout for a larger branch office, which is distinguished from a small branch office by having a local domain controller on site. A single IAS server is deployed to the remote office. Although the IAS server is depicted as a separate server, you could run this service on the domain controller.

Note: If the WAN link to the head office is reliable (that is, there are redundant network links) and not overly congested, the large branch office can use the head office RADIUS services rather than its own. This option is discussed further in Chapter 5, "Designing the RADIUS Infrastructure for Wireless LAN Security."

All other services (for example, the CAs) are supplied from the head office.

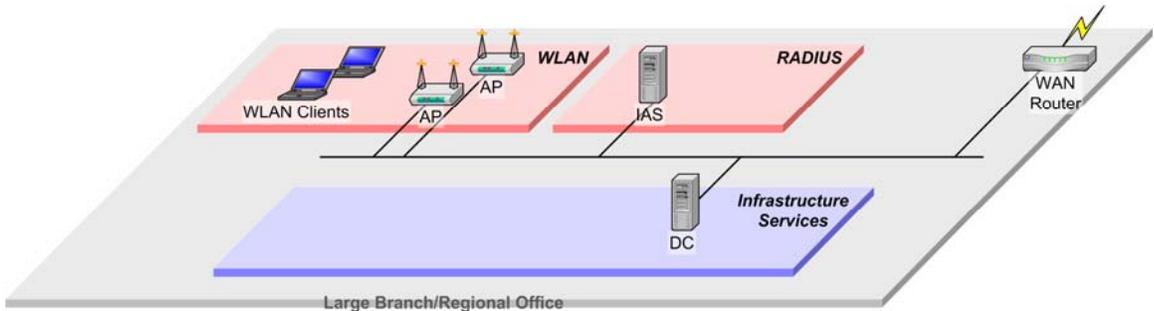


Figure 3.7
The physical layout for a large branch office

Small Branch Office

The small branch office may have some IT infrastructure—for example, a file server and printer — but it will typically not have any authentication infrastructure. Some organizations believe that these offices do not require or justify any WLAN services. Other organizations, using temporary offices, find that not needing to lay and manage network cables is an attractive option.

If WLAN services are needed in a small office that lacks a local domain controller, the local wireless APs will rely on the IAS server and domain authentication infrastructure at headquarters. The main problem with this approach is that all WLAN connectivity is lost if the WAN link to the head office fails. Although there is no easy solution to this scenario, you can address this weakness (at a cost) by providing WAN redundancy or by deploying local domain controllers.

If WAN resilience or locating local domain controllers in the small branch offices of your organization is too expensive, an alternative is to deploy isolated wireless APs using the WPA pre-shared key (PSK) mode. All Wi-Fi certified wireless APs now support WPA. Although this is much more secure than static WEP, there is additional management overhead associated with this option.

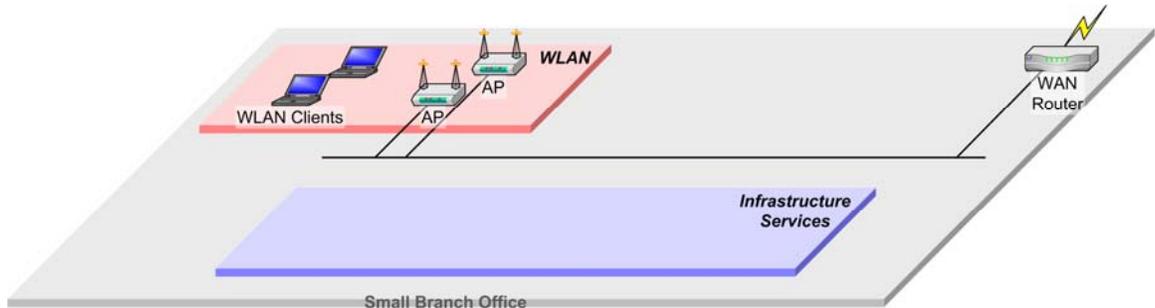


Figure 3.8

The physical layout for a small branch office

Scaling Strategy

One of the key design criteria is to ensure that the design can scale. The solution has to support a wide range of implementation sizes at a cost that is appropriate for each one. For example, an implementation for 500 users should cost proportionally less than one for 5,000 users. The complexity of implementing and managing the solution must also be realistic across this range of organizations.

Large Organization

The following figure illustrates how the design would scale to accommodate large numbers of users in a head office and larger regional offices. It is possible that the IAS servers would also service other network applications, such as VPN. For more information about this subject, see the "Extending the Design" section later in this chapter. This consideration could also potentially influence the precise layout and number of the servers. The extra IAS servers are displayed in the following figure for illustration only.

The additional servers required for the scaled version of the solution are shaded.

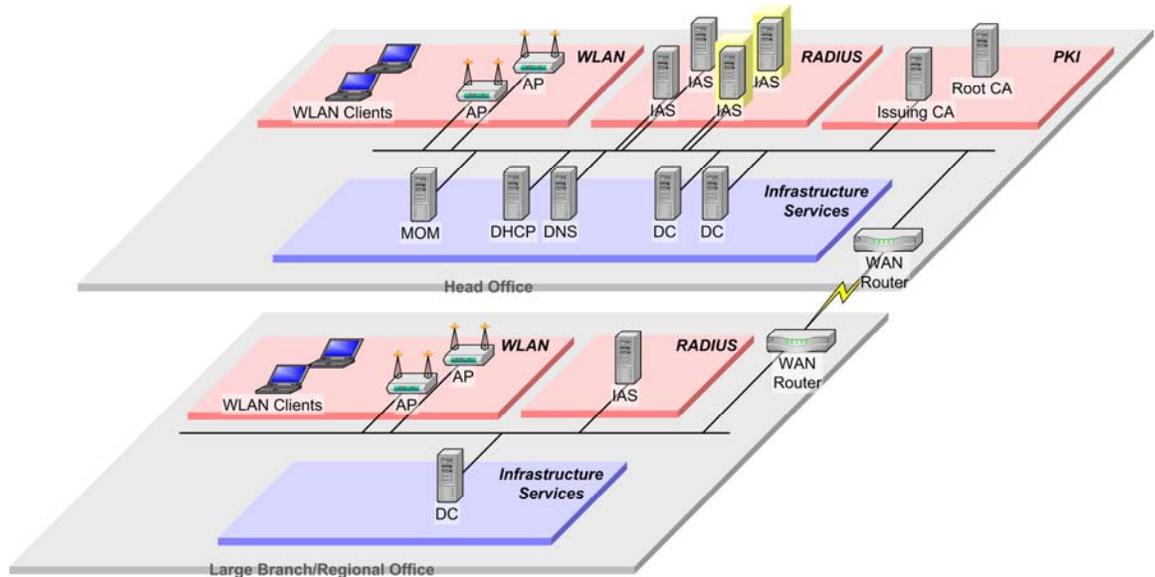


Figure 3.9

The solution scaled for a large organization

Small Organization

At the other of the spectrum, the solution can be implemented with a relatively modest amount of new hardware and software. This is mainly achieved by running the IAS service on existing domain controllers. This configuration has been extensively tested by the IAS product group at Microsoft and is recommended for many scenarios. The following figure illustrates this version of the design.

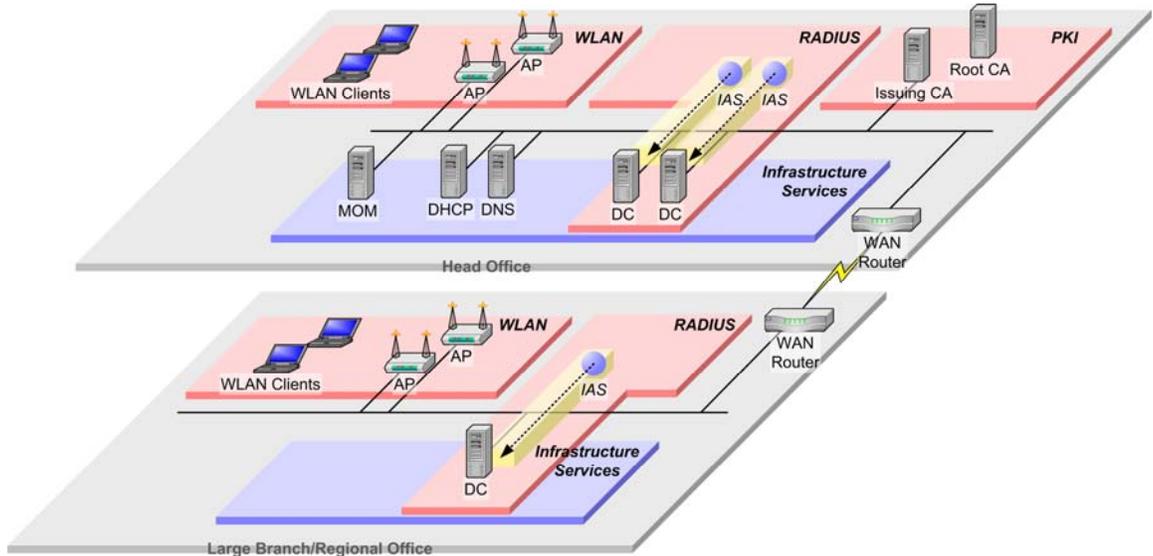


Figure 3.10

The solution scaled for a small organization

The RADIUS component is still displayed in the figure as logically separate (in order to match the layout of the previous figure for comparison) but the component is actually implemented as a service on the pre-existing domain controllers. The only servers required in this version of the solution are the CAs that reside in the PKI area of the solution design.

Extending the Design

Another key design criterion of the solution is the reusability of the components in future applications. You can reuse both the RADIUS component and the PKI component to provide authentication and other security services for a variety of applications.

Other Network Access Services

This solution's RADIUS design can provide authentication, authorization, and accounting services for other network access servers, such as 802.1X wired network authentication, and VPN and remote access authentication.

802.1X Wired Network Authentication

The simplest application, which requires no modification of the basic WLAN RADIUS design, is 802.1X wired authentication. Organizations that have a widely distributed wired network infrastructure may find it difficult to control unauthorized use of the corporate network. For example, it is often difficult to prevent visitors plugging in laptops or employees adding unauthorized computers to the network. Some parts of the network, such as data centers, may be designated high security zones. Only authorized devices should be allowed in these high security zones — even to the exclusion of employees with corporate computers.

The following figure displays how a wired network access solution integrates with the design: The bold-edged area represents the 802.1X wired components, while the other areas of the layout contain the relevant services displayed in the previous design figure.

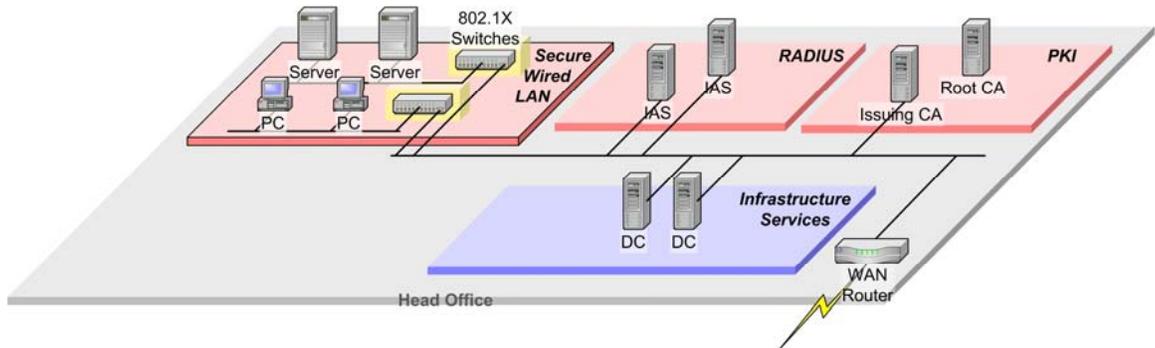


Figure 3.11

Using 802.1X wired authentication

Networks that use 802.1X switches play a role that is identical to the wireless APs role in the core solution. Moreover, these networks can use the same RADIUS infrastructure to authenticate clients and selectively authorize access to the appropriate network segment. This version of the solution includes the obvious advantages of centralizing the account management in the corporate directory while leaving the network access policies under the control of the network security administrator.

VPN and Remote Dial-Up Authentication

Another network access service that could use the RADIUS components is VPN and remote dial-up. Particularly in larger organizations, it is likely that some additions would need to be made to the design as it stands, such as the addition of RADIUS proxies. The following figure displays how the extended solution might look.

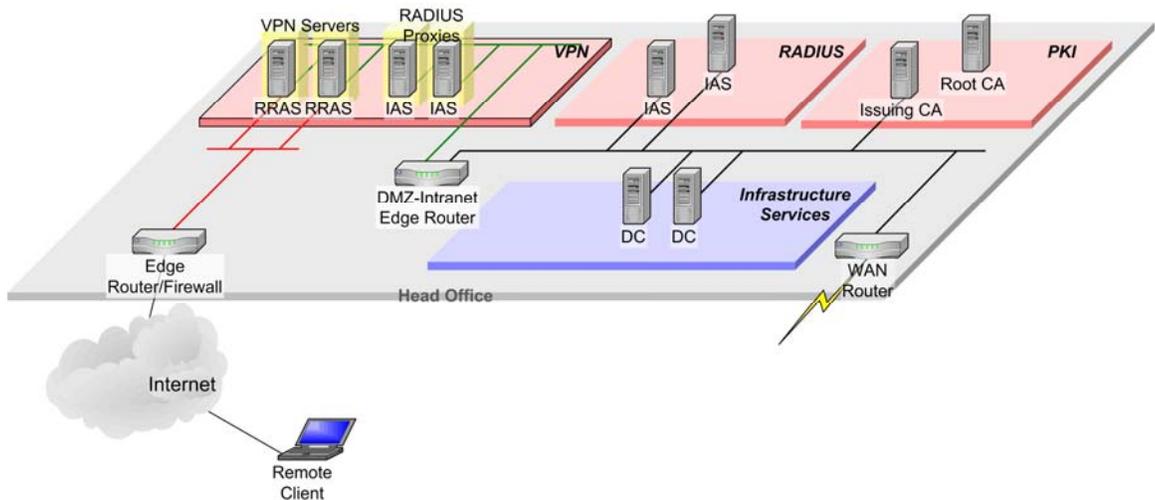


Figure 3.12

Extending the RADIUS component to support a VPN

The VPN servers in this solution play the same NAS role as the wireless APs in the core design: They pass the clients' authentication requests to the RADIUS infrastructure. Although it is possible to directly pass the RADIUS requests to the internal IAS servers, it is more secure to use a RADIUS proxy layer that forwards the requests to the internal IAS servers.

This solution combines the advantages of using existing infrastructure and centralizing account management, while leaving the access policy control under the supervision of the network security administrator. Further enhancements, such as mandating smart card-based user authentication, add to the overall security of the solution. Microsoft uses a very similar configuration to allow its internal staff to securely connect to the corporate network.

Dial-up remote access works in a similar way by using the dial-up server capability of the Windows Routing and Remote Access service instead of the VPN functionality.

Using RADIUS (specifically IAS) in this scenario offers another advantage: the ability to use *quarantining* policies. This takes advantage of the Routing and Remote Access service in Microsoft Windows Server 2003, and Connection Manager (the Windows-enhanced remote access client) to allow or deny access based on the security state of the client computer. With this configuration, IAS can verify the client meets certain requirements when it connects to the network. For example, this procedure can check to ensure that the client has up-to-date antivirus software or is running a corporate-approved operating system build. If the client fails either of these checks, the RADIUS server denies it access to the network. In this way, even a properly authenticated user and computer may be denied access if they present a possible security threat to the company network.

PKI Applications

Because the solution criteria for reusability and extensibility are important, the PKI component was designed in the knowledge that it may be used in the future for a variety of different security applications. As the next chapter will discuss, the PKI design is therefore a blended strategy; minimizing cost and complexity as part of a secure wireless solution, while also maintaining enough flexibility for you to use it as a basis for other applications in the future.

The following figure illustrates a few of the applications that the PKI component could support in addition to the secure wireless application. Some are relatively simple applications that can use the PKI developed in this solution with little or no modification to the core design. Others, such as secure e-mail and smart card logon, are more complex, and will almost certainly require more careful consideration and some extension of the PKI design.

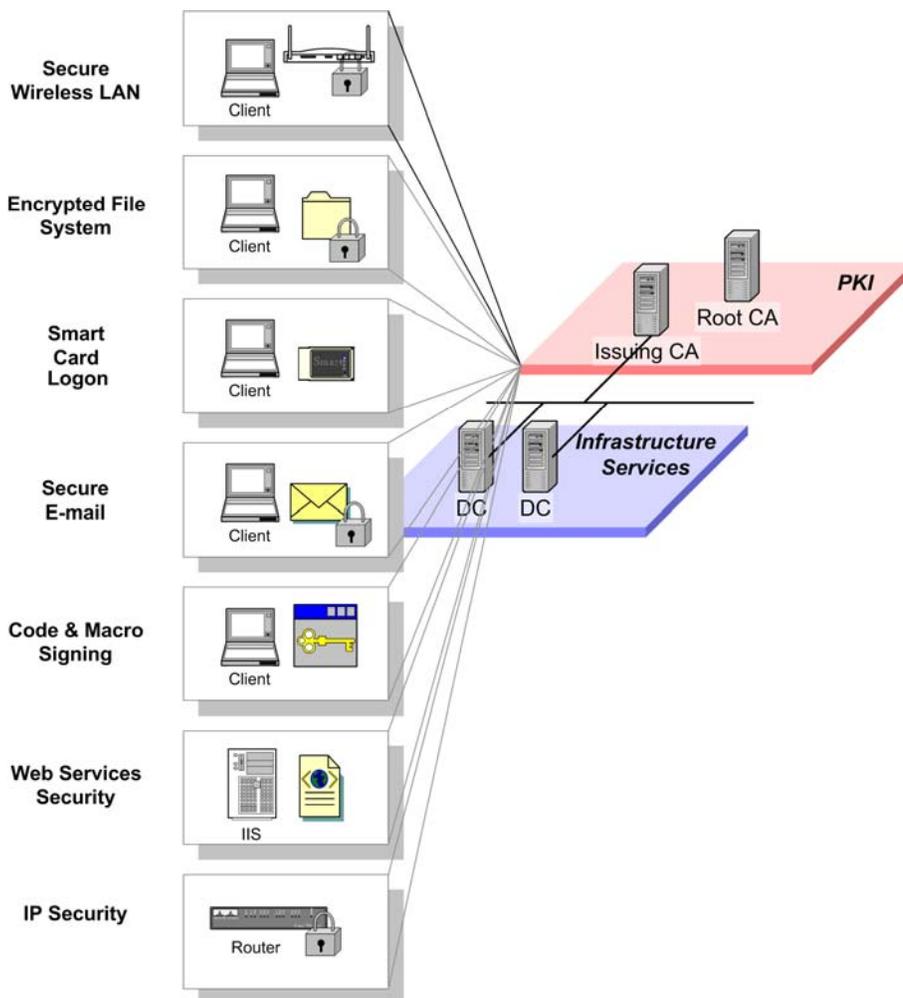


Figure 3.13
PKI applications

Design Criteria Re-Evaluated

Before closing this chapter, it is worth re-examining the list of design criteria for the solution to examine how well the proposed design now meets the goals set earlier. This evaluation is summarized in the following list. However, many of these items are only fully addressed in the detailed design chapters that follow this one.

- **Security.** The solution design includes robust authentication, authorization, and access control. Strong (128 bit) encryption is a function of the network hardware and is supported on most currently available devices. Secure management of the encryption keys is provided by a combination of the Microsoft 802.1X client, the 802.1X-enabled wireless AP and the wireless network cards, and the RADIUS server.

Achieving resiliency in the face of DoS attacks remains an area where there is still work to do—current industry standards (until the advent of 802.11i) are still vulnerable to a variety of DoS attacks.

- **Scalability.** The basic design accommodates a wide range of organizations in a cost-effective manner from a few hundred to many thousands of users. The design is also flexible with regard to geographic and network layout. Small offices without a local domain controller are dependent on WAN reliability or a lower grade security solution.
- **Component reuse (use of existing infrastructure).** The design uses Active Directory and many existing network services, such as Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS).
- **Component reuse by future applications.** The RADIUS design, implemented using IAS, can be used by or easily extended to support other network access applications (such as VPN, 802.1X wired network access, and remote access dial-up). The PKI is also capable of supporting simple public key applications, such as EFS, and provides the environment to work with more complex applications that can perform such things as smart card logon.
This item also meets the design criterion — **Extensibility.**
- **Availability.** The solution design is resilient to a single component or network link failure at the head office, and for all outlying offices where a RADIUS server can be deployed. Small offices without a local RADIUS server are vulnerable to a WAN failure.
- **Manageability.** The ability to manage the solution is not apparent from the design, but this requirement is accounted for in the design of the operational framework
- **IT organization structure.** Some level of specialization with WLANs in the organization's IT department is essential for deploying and managing a solution of this type.
- **Standards compliance.** The solution adheres to current official and industry standards. This is most relevant in the area of WLAN security where the solution is based on the 802.1X protocol, EAP-TLS, and 128-bit dynamic WEP or WPA. Microsoft recently announced product support for WPA for Windows XP, approving the highest available standards of native WLAN security. The design will support either WPA or dynamic WEP.

Summary

This chapter examined the conceptual design of a secure wireless LAN network solution based on using the 802.1X protocol and EAP–TLS . The key components of the solution were explained at an architectural level. An outline of the target organization for this solution was then described, together with the design criteria used to engineer the solution.

The design criteria were used to translate the conceptual solution into a logical solution design. This included examining implementation options to scale the solution for organizations of different sizes with different requirements, and how to extend the basic design to provide support for other network access and security applications. Finally, the main design criteria were reviewed against the features of the proposed design. This criteria review serves as an entry point to the remainder of the Planning Guide chapters.

The next three chapters of the guide detail the design of each major architectural components of the solution: the PKI, the RADIUS infrastructure, and the WLAN security design.

4

Designing the Public Key Infrastructure

Introduction

The previous chapter described a logical design for a secure wireless solution that depends on a Public Key Infrastructure (PKI). This chapter defines the process of designing a PKI based on Microsoft® Windows® 2003 Certificate Services for this solution. In order to keep deployment and management costs down, the solution design is relatively simple and well suited to issuing certificates for secure wireless clients, and for the wireless local area network (WLAN) infrastructure.

However, while the primary objective is to design a PKI that will support secure WLANs, keep in mind that a PKI also can form an important part of your organization's overall security infrastructure—one that a variety of other applications in your environment can use in the future. In order to protect your investment in this infrastructure, the solution design is extensible. This means that while the design may not be suited to issuing all types of certificates, it will allow you to add additional functionality and capacity in the future to meet a broader set of security requirements than those addressed here.

This chapter has three main objectives. The first is to discuss the solution design decisions and the reasoning behind them. The second is to give you some background planning information to help you decide whether those decisions are right for your PKI. The third is to point to ways in which you can extend the basic solution to meet security needs that are outside the scope of this solution.

When phrases like “This solution uses option...” or “This design uses...” appear in this chapter, they refer to the decisions made as part of the solution design that are implemented in the Build and Operations Guide chapters for the solution.

When phrases such as “You should decide this...” appear, they indicate items where you must decide something based on your own requirements. This will mostly occur where the text is discussing how you might extend the solution to meet your organization's broader security needs. For this reason, some topics feature a more detailed discussion in order to help you understand the implications of the step and to prevent you from having to refer to other documents.

Chapter Prerequisites

You should have a good understanding of the general principles and terminology of PKI. If you are new to the technology, read some of the articles referred to in the “More Information” section at the end of this chapter.

Before you continue with this chapter, you should familiarize yourself with the “Designing a Public Key Infrastructure” chapter of the *Microsoft Windows Server™ 2003 Deployment Kit*. See the “More Information” section at the end of this chapter to learn where to find this information. This chapter follows the structure of the “Designing a Public Key Infrastructure” chapter in the deployment kit to make it easy for you to refer to the relevant background information and more detailed discussion contained there.

Links to additional detailed information about how to plan and design a Windows Server 2003 PKI are also available in the “More Information” section.

Chapter Overview

Planning and deploying a PKI that meets your organization's current and future needs is not a trivial task. Usually, a PKI is not intended to provide a solution to a single, isolated security problem. Instead, an organization deploys a PKI to address a number of internal security requirements, as well as business security requirements to work with external customers or business partners.

The following flowchart illustrates the chapter structure.

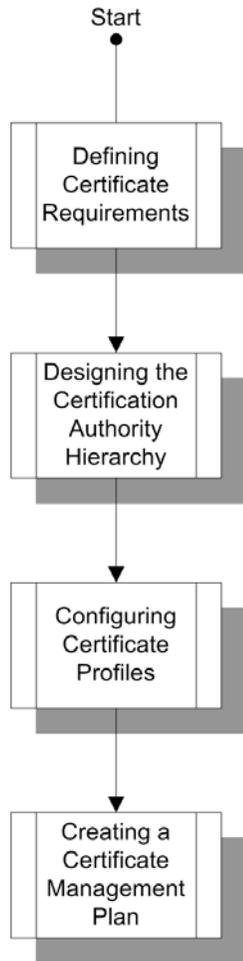


Figure 4.1
Chapter structure for planning Certificate Services

These four main steps are:

- **Defining Certificate Requirements.** This step involves defining the security problems that you are trying to solve. This is based on the specific applications and users that need enhanced security, where these users are located, and the degree of enhanced security that is required. You cannot begin to create your PKI until you have defined your security and business requirements.
- **Designing the Certification Authority Hierarchy.** Based on a variety of factors, you must create an infrastructure of certification authorities (CAs). This step involves defining a trust model, determining how many CAs you need, how you will manage them, and how you can extend your PKI by introducing additional CAs or establishing trusts with other organizations. In addition, this step addresses how the PKI integrates with other technologies in your IT infrastructure, such as the Active Directory® directory service, and Microsoft Internet Information Services (IIS).
- **Configuring Certificate Profiles.** This step includes deciding what types of certificates to use, how strong the encryption keys associated with those certificates must be, what is the valid lifetime of the certificates, and whether they are renewable.
- **Creating a Certificate Management Plan.** This step defines how certificates are issued to end users, how certificate requests are processed, and how certificate revocation lists (CRLs) are managed and distributed.

Defining Certificate Requirements

This section defines the purposes for which the PKI issues certificates and the security requirements for each purpose.

Creating a Certificate Practices Statement

As you design your PKI, you should record decisions about how certificates will be issued and used in your organization. These decisions are referred to as certificate policies, and the documents recording them are known as certificate policy statements and certificate practices statements.

In formal terms, a *certificate policy* (CP) is a set of rules under which the PKI operates. It records, for example, the applicability of a certificate to a particular group of clients or applications with common security requirements. A *certificate practices statement* (CPS) is a statement of the practices that an organization uses to manage the certificates that it issues. It describes how an organization's certificate policy is interpreted in the context of the organization's system architecture and operating procedures. A CP is an organization-wide document. However, a CPS is specific to a CA (although CAs can have a common CPS if they perform the same job—for example, if you are splitting the CA load across multiple servers for performance or resilience reasons).

For some organizations and certificate uses, the CP and CPSs are considered legal documents or legal disclaimers. These typically require specialized legal advice that is outside the scope of this chapter. However, there is no strict requirement for you to produce either of these documents as part of your PKI. Unless you have specific legal or commercial reasons for doing so, you may want to avoid the time and expense of producing and maintaining formal certificate policy and practices statements.

Although you may not need a formal CP or CPS, you should still document your certificate policies and operational practices. Certificate policies should become part of your organization's overall security policy, and the operational practices should become part of your security management procedures. You might refer to this as an informal CPS.

Based on the intended uses for your PKI, you should decide whether or not to produce a formal policy statement and CPS. If you require a formal CPS, you will probably need to publish it and reference it in your CA certificates. Although no guidance on writing a formal CPS is included in this solution, instructions on how to publish one are included in Chapter 7, "Implementing a Public Key Infrastructure." You do not normally need to publish informal CPSs.

In the remainder of this chapter, there are frequent references to documenting decisions in your CPS. These instructions apply equally to a formal and an informal CPS.

You will find additional sources on producing a CPS in the "More Information" section at the end of this chapter.

Identifying Certificate Applications

The first step in the PKI design process is to identify the list of applications that will make use of certificates. For each application you should document the types and approximate number of certificates required for each application. You do not need to specify any certificate details at this stage; you should just provide a brief description.

The secure wireless solution requires certificates for wireless clients and for the Windows Remote Authentication Dial-In User Service (RADIUS) servers. The Microsoft RADIUS server is a component of Windows Server, called the Internet Authentication Service (IAS).

The required certificate types are displayed in the following table. Although not strictly required for this solution, the PKI will also issue certificates to domain controllers (this is the default when a Windows 2003 Enterprise CA is installed into the forest).

Table 4.1: Certificate Requirements for the Secure Wireless Solution

Application	Certificate type	Number of certificates
Secure WLAN	Client authentication certificates for users.	All users who require WLAN access.
	Client authentication certificates for computers.	All wireless LAN computers.
	Server authentication certificates for the IAS servers.	All IAS servers.
Active Directory	Domain controller authentication.	All domain controllers in the forest.

In the future, you can extend the PKI to issue certificates for the applications displayed in the following tables.

Table 4.2: Potential Future Certificate Requirements

Application	Certificate type	Number of certificates
Client access virtual private network (VPN)	Computer client authentication (IPsec)	All remote VPN clients
Branch-to-branch VPN	VPN server authentication (IPsec)	All VPN routers
IP Security (IPsec)	Computer client authentication	All client and server computers requiring IPsec.
Web security	Authentication of users to intranet Web applications.	All users
	Intranet Web server	Secure intranet Web servers
Encrypting File System (EFS)	EFS user	All users
	EFS Data Recovery	Recovery agents
Secure e-mail	Secure/Multipurpose Internet Mail Extensions (S/MIME) signing and encryption	All e-mail users
	Key recovery	Recovery agents
Smart cards	Smart card login	Domain users
Code signing	Internal code and macro signing	Code release manager

Defining Certificate Clients

For the applications listed in the preceding section, you should define the clients that will use the certificates. The term “client,” in this context, means any person, software process, or device that uses the certificates issued by the PKI. For example, clients include users, servers, workstations, and network devices. In order to understand how the issued certificates will be used, you must consider two major categories of clients: the certificate subject (or *end entity*), and other certificate users.

End entities are clients that have a certificate issued to them by the PKI. The certificate will have one or more entries in its **Subject** or **Subject Alternative Name** fields that identify the client (for example, host name, e-mail address, or directory distinguished name) as the owner of that certificate. The other category of certificate users are clients that may need to verify the certificates of end entities or look them up in a directory, for example, but who do not necessarily have certificates issued to them by the PKI.

A common example helps clarify the distinction: An internet user buying something on a secure Web site will be a user of the Web site's Secure Sockets Layer (SSL) certificate. The Web site, though, is a certificate end entity; its identity—`www.woodgrovebank.com`—is encoded into the certificate **Subject** field. Only the certificate subject has access to use the certificate private key—other users of the certificate do not. Note: Certificate *subjects* are nearly always also certificate *users* of their own and, usually, others' certificates.

Note: “End entity” is the technically correct term, but the friendlier term, “certificate subject,” is used instead throughout most of this chapter.

For both certificate subjects and certificate users, you should categorize each client type by answering the following questions:

- Is the client a person, a computer or device, or a software process?
- On what platforms (operating system version) will the certificates be used?
- What is the network location of the client? For example, is it connected to the internal LAN, in a partner organization or on the Internet?
- Is the client a domain member? If so, is it in a different domain or a different forest from the CA? Is it an untrusted domain?
- What type of operations does the client need to perform? For example, enroll certificates, sign with certificate, verify certificate trust, look up certificates in a directory, and check certificate revocation status.

This categorization will impact many design decisions, such as how the certificate is issued, the level of trust that you can place in a given certificate, and how certificate revocation information is published.

For this solution, the client categories are detailed in the following tables.

Table 4.3: Certificate Subject (End Entity) Categories

Certificate	Client type	Platform	Location	Domain	Certificate operations
Wireless Client authentication	User	Windows XP	Internal network	Domain member	–enroll –authenticate
Wireless Client authentication	Computer	Windows XP	Internal network	Domain member	–enroll –authenticate
IAS Server authentication	Computer	Microsoft Windows Server™ 2003	Internal network	Domain member	–enroll –authenticate –secure channel

In this application, the users of the certificates will be the same set of clients but with the roles reversed. For example, the IAS server becomes the user of the client certificates and needs to verify them. Verification usually includes verifying that the certificate chains to a trusted root CA, and that the signature supplied by the client matches the public key in the client's certificate. The certificate may also be subject to a revocation check. For a detailed description of this topic, see the paper, *Troubleshooting Certificate Status and Revocation*. A full reference is located in the “More Information” section at the end of this chapter.

Table 4.4: Certificate User Categories

Certificate	Client type	Platform	Location	Domain	Certificate operations
Wireless Client authentication	–Computer	Windows Server 2003	Internal network	Domain member	–verify –check revocation
Wireless Client authentication	–Computer	Windows Server 2003	Internal network	Domain member	–verify –check revocation
IAS Server authentication	–User –Computer	Windows XP	Internal network	Domain member	–verify

From the previous tables you can identify the platforms and the kinds of operations that you need to support. Although not the case in the WLAN scenario, for other applications in your environment, you may need to support certificate lookup or verification by clients on the internet or enrollment from non–Windows platforms. Because you must decide many of these items early in the design process, it is important to think about your probable future certificate requirements.

This solution makes the following assumptions about future requirements:

- Certificate verification from non–Windows clients will likely be required.
- Certificate verification from the Internet may be required.
- Support for both certificate subjects and certificate users on platforms other than Windows XP and Windows Server 2003 will be required.

Although the design does not necessarily meet all of these requirements at the moment, it will be easy to accommodate them in the design.

Defining Certificate Security Requirements

The security of a certificate is also known as its assurance level. It can be thought of as a measure of the strength that binds the subject of the certificate to the certificate itself. It reflects how confident you can be that the person (or device) using the certificate is really the same as the subject named in the certificate. The assurance level is a measure of two main things:

- The rigor of the registration and certificate enrollment process—for example, did the person need to show up in person and present photo ID to get his or her certificate, or was possession of an e-mail address sufficient?
- The way in which the private key is stored—the more difficult it is to copy or otherwise compromise the key, the stronger your assurance that it is still in the unique possession of the original owner—the certificate subject.

The two are closely linked since there is no reason to invest in expensive private key protection measures if you were never really sure of the identity of the private key owner. Similarly, an arduous registration process involving extensive background checks and DNA fingerprinting is of little value if the private key is subsequently stored in a way that is not secure.

Achieving a higher assurance for a certificate costs money, though, and is frequently not necessary for many certificate uses. If you do not want any stronger assurance from a certificate than that it belongs to an authorized domain user, then domain credentials are completely acceptable as the registration evidence to enroll a certificate.

You should document the meaning of the assurance levels that you use in your certificate policy and practices statements.

For this solution, the following table defines three assurance levels.

Table 4.5: Certificate Assurance Levels

Level	Registration requirements	Key storage requirements
Standard (Low)	Automatic approval dependent on domain or other password based identification.	Software keys
Medium	Certificate Manager approval, visual ID check (smart cards), or enrollment officer signature.	Software keys or hardware tamper-proof token (smart card or USB token).
High	Nominated enrollment officer signature and Certificate Manager approval.	Hardware tamper-proof token (smart card or USB token).

There is some overlap between these categories. They are not strict technical divisions; they are really policy divisions. In your organization the boundaries between them will be based on policy decisions that you make about how you want your certificates to be treated. In general, you can expect high assurance certificates to be somewhat rare, whereas standard assurance certificates tend to be very common.

Important: This chapter uses the terms “standard value” certificates and “standard assurance” certificates, instead of “low value” and “low assurance.” Because these latter terms have a negative connotation, “standard” better reflects the intended meaning.

You can further refine the assurance categories defined in the previous table by dividing each into different subject types. Common categories for these are:

- Computer—this is really any computer or device within your organization.
- Internal user—this represents full-time employees or staff who you consider equivalent (for example, contract staff).
- External user—this represents any other entity who exists outside of your organization with whom you have some type of business or legal relationship (for example, business partners and customers).

The reason for this distinction is that these different subject types usually have quite different certificate policies applied to them; that is, the conditions under which a certificate is issued, revoked, or renewed. Even if you have no certificate plans for a given category, you may want to document the certificate policies that would apply to that category so that your policies and CPS are properly prepared. The following table describes the results of combining the assurance levels and subject categories.

Table 4.6: Certificate Security Categories

Certificate security category	Example characteristics of security category	Example certificate types
Computer certificates		
Standard assurance computer certificates	<ul style="list-style-type: none"> –Automatic approval based on computer domain credentials. –Yearly renewal. 	<ul style="list-style-type: none"> –WLAN computer –IPsec
Medium assurance computer certificates	<ul style="list-style-type: none"> –Certificate manager approval required. –Key storage in software. –Yearly renewal. 	<ul style="list-style-type: none"> –Web server –IAS Server authentication
High assurance computer certificates	<ul style="list-style-type: none"> –Certificate manager approval. –Key storage in hardware security module (HSM). 	<ul style="list-style-type: none"> –Certificate Authority –Secure Time Service –Registration Authority
Internal user certificates		
Standard assurance internal user certificates	<ul style="list-style-type: none"> –Automatic approval based on user domain credentials. –Yearly renewal. 	EFS user
Medium assurance internal user certificates	<ul style="list-style-type: none"> –Certificate manager or enrollment officer approval required. –Key storage on smart card or software. –Yearly renewal. 	<ul style="list-style-type: none"> –Secure e-mail –Low-medium value financial authorization –Smart card logon –Internal code signing –Data recovery agent –Key recovery agent
High assurance internal user certificates	<ul style="list-style-type: none"> –Physical ID verification of certificate subject required. –Certificate manager approval required. –Enrollment office signature required on request. –Key storage on smart card. –Six month renewal. 	<ul style="list-style-type: none"> –High value financial authorization –Commercial code signing

(continued)

External (user) certificates		
Standard assurance external certificates	–Automatic approval based on pre-assigned password. –Yearly renewal.	Client authentication (authentication to internet Web site)
Medium assurance external certificates	–Certificate manager approval required. –Key storage on smart card. –Six month renewal.	Business-to-business (B2B) financial authorization
High assurance external certificates	–Physical ID verification of certificate subject required. –Certificate manager approval required. –Enrollment office signature required on request. –Key storage on smart card. –Six month renewal.	Very high value B2B transaction

Note: If you have no use for a particular categorization, you do not need to create one. You may want to use a simpler or more complex classification system, and not every combination need result in certificate type that you are going to issue.

There are no technical reasons why these different certificate subject types cannot all be treated in the same way. However, you will typically define different security policies for different subject types; for example, internal employees will be treated differently than staff in other organizations. The different certificate policies (and their embodiment in different CPSs) may affect your decisions on how to structure your CAs to issue these different certificate types. This will be covered later in the chapter.

You should also consider whether the same administrator will ultimately be responsible for certificates issued to these three categories of certificate users (or “end entities” in PKI terminology). In many organizations, the person who can certify that a computer is a legitimate domain member is not the same person who can certify the identity of a partner company. You should document these responsibility relationships in your CPS.

Defining Application Certificate Security

The certificate security categories defined in the previous section can be used to classify the certificate types for the design. This is listed in the following table.

Table 4.7: Certificate Security Requirements

Certificate type	Security category	Platform	Logical location	Approval	Key size	Validity period
Client Authentication — User	Standard assurance computer certificates	–Windows XP –Windows Server 2003	Internal	Automatic (domain authentication)	Medium	Medium
Client Authentication — Computer	Standard assurance user certificates	–Windows XP –Windows Server 2003	Internal	Automatic (domain authentication)	Medium	Medium
IAS Server Authentication	Medium assurance computer certificates	–Windows XP –Windows Server 2003	Internal	Manual	Medium	Medium

These broad requirements will be refined into specific certificate profiles in the “Configuring Certificate Profiles” section later in the chapter.

Combining Certificate Purposes

It is possible to combine a number of application functions (or usages) onto a single certificate so that you can use one certificate to sign e-mail, log on to the network, and grant access to an application. Combining usages will result in lower management and storage overhead on the certificate and directory servers.

However, there are disadvantages to multi-purpose certificates. For example, different applications may require a different approval process for the certificate. Most reasons for using multiple certificates are technical, but the chief one is usually that different applications require different certificate security levels; that is, different assurance levels bind the certificate to the certificate subject. This can include differences in any or all of the following:

- Certificate approval process
- Key length
- Key storage mechanism
- Certificate lifetime

Because of this, the strategy of combining certificate usages with the same security level, is usually the best one. The client authentication certificate type used in this solution includes usages for other standard applications such as IPsec and computer authentication. As you define requirements for other certificate usages, you can include them and then re-issue the certificates, which will require forced renewal that you can initiate from the certificate template definition.

However, the IAS server certificates are considered medium assurance certificates. The threat posed by an unauthorized server certificate is much greater than that posed by an illegitimate client certificate. For this reason, it is sensible to treat server certificates with more care, and Microsoft recommends not combining them with usages for standard assurance applications.

Designing the Certification Authority Hierarchy

To support your organization's certificate-based applications, you must establish a framework of linked CAs that is responsible for issuing, validating, renewing, and revoking certificates as needed. The CAs in turn rely on an underlying IT infrastructure for things such as certificate subject authentication, certificate publishing, and certificate revocation information publishing.

The goal in establishing a CA infrastructure is to provide reliable service to users, manageability for administrators, and flexibility to meet both current and future needs, while maintaining an optimum level of security for the organization.

Selecting a Trust Model

The first step in designing your CA infrastructure is to determine the trust model that is the most appropriate for your requirements. The two basic models are the Hierarchical Trust and the Network Trust, although it is possible to combine elements of both models into a Hybrid trust model. For further discussion of these models, see the "Designing a Public Key Infrastructure" chapter of the *Windows Server 2003 Deployment Kit* in the "More Information" section.

This solution in this guide uses the Hierarchical Trust model. The reasons for this are:

- You can treat offline CAs with a greater level of security than online CAs. One or more layers of offline CAs increases the overall level of trust possible in the issued certificates.
- Hierarchies can work more easily without the presence of a directory; this is important for supporting external clients who have no access to your internal directory. Network trusts usually require directories so that users can look up CA cross-certificates in order to build trust chains. The trust chain in hierarchies is always explicit.
- There are fewer trust anchors to maintain and distribute to clients—you only must distribute the root CA certificate to your certificate users.
- Even with a rooted hierarchy, the option always exists to include multiple trust anchors (or roots) in the future by cross-certifying with other hierarchies. This means that the design can accommodate things like organization mergers, and devolving control of certificates for special purposes to departments.

A single root is adequate for the proposed solution.

Third Party vs. Internal Root

It is possible to use an internal root as the trust anchor for the PKI or to use the services of a commercial CA for this. Using a third-party root implies that your issuing CAs are certified by the commercial root CA (usually via one or more intermediate CAs). Therefore, all of your issued certificates ultimately have their trust anchors at this external root CA.

Note: Although this is not considered explicitly in this guidance, it is possible to outsource all of your organization's certificate requirements to a commercial CA. You can use either an on-site managed service or obtain the certificates directly from the certificate provider. This is often not a financially viable alternative except for small organizations or for very restricted certificate usage. This decision is completely unrelated to the question of whether to use an internal versus a third-party root — although the two things often confused.

There are a number of advantages to using a commercial root for your internally issued certificates:

- A commercial root gives external parties (for example, customers visiting your secure Web site or partner organization receiving a signed e-mail) a greater degree of confidence when conducting secure transactions with the organization. They will typically already trust the third-party root CA, and so they will not have to decide whether they trust your certificates.
- A commercial root allows the organization to take advantage of the expertise of a professional service provider, including the provider's understanding of the technical, legal, and business issues associated with certificate use. (Although, unless the certificate provider is issuing all of your certificates, you still have responsibility for the way that your certificates are issued and used, and you should document this in your CPS.)

However, there are a number of disadvantages to this approach:

- It will typically involve a high per-certificate cost.
- The certificate provider may require stringent security and audit measures for all CAs subordinated by the commercial root CA.
- Internal users and devices must have access to the third-party CAs' CRLs published on the Internet.
- Some applications may require specific parameters or extensions in the root and intermediate CA certificates (for example, Microsoft Windows smart card logon), which may not be available from the certificate provider.
- The commercial agreement between your organization and the certificate provider may restrict the type of certificates that you are able to issue from subordinate CAs. For example, Web server certificates may not be allowed.
- Trusting a commercial root CA may be too wide a scope of trust for the security needs of your organization. You may have to introduce special checks or an extra layer of internal CAs to distinguish between certificates issued by your organization and those issued by another organization also subordinated to the same root.

Despite these disadvantages, if you need to issue significant numbers of certificates that users outside your organization will trust, you should consider subordinating at least a part of your PKI beneath a commercial root (although this may require creation of two separate hierarchies).

For the majority of certificates used in the organization, this solution uses a hierarchy based on an internal root CA. This approach has the following advantages:

- It allows the organization to maintain direct control over the central trust anchor—the root CA—and the security policies that govern the issuance and use of certificates issued by it.
- Large numbers of certificates can be issued from the internal PKI at relatively low cost.
- There are no restrictions on the types of certificates that you can issue.
- There is no ambiguity between the trust in internal certificates and external certificates.
- You can publish CRLs and Authority Information Accesses (AIA) information internally or externally as required.

If you cannot easily manage the trusted roots of your certificate users you should consider using an external root. This solution proposes using third-party certificates and an external root for the following services:

- Internet Web server
- Commercial code signing
- Commercial document signing
- Externally trusted secure e-mail

Defining External Certificate Trusts

The previous section touched on trust in the certificate infrastructure of other organizations. You must consider this subject more broadly to determine how trust in certificates is controlled in your organization. The word “trust” in this context has three significant conditions:

- The person or entity in which the trust is placed—whom do you trust?
- The operations or activities that you trust the party to perform —what do you trust them to do?
- The time period in which you want to maintain that trust—how long will you trust them?

For certificates, the “whom” is the certificate issuing authority, and the “what” is the usages and other certificate characteristics that you may want to control. The “how long” is either defined by the validity period of the root CA certificate or, in some cases, the validity period of a special cross-trust certificate that you create.

It is likely that you will need to change the default trust relationships that your organization has with external parties when you establish a new business relationship with another organization or want to enable some function for your users (for example, trusting Web certificates to allow secure HTTP sessions). For example, some things that you might want to do are:

- Distribute the CA certificate of a partner organization (or a new commercial certificate provider) so that some or all of your users trust the partner or commercial CA certificates.
- Distribute the CA certificate of a special purpose CA or PKI within your organization that you do not want the whole organization to trust.
- Replace the pre-existing commercial roots in your clients' root store so that you can restrict the usages of your trust certificates. For example, you might decide that you only trust a given commercial root for e-mail and secure Web server certificates, but not, say, for smart card logon certificates.

There are a number of ways that you can achieve these goals:

- Create qualified subordination relationships between your internal root and the CA certificate to be trusted (also known as cross-certification). This involves one of your internal CAs re-signing the external CA certificate. This effectively adds the external CA into your internal PKI as a trusted subordinate of the signing CA. You can place restrictions on the type of certificate, limiting precisely the certificate usages and policies, types of subject names, or issuance policies that you will trust.

Important: The subject of qualified subordination or cross-certification is a complex one and the most difficult of the methods to successfully implement. See the technical paper, *Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003* in the “More Information” section at the end of the chapter.

- Create a Certificate Trust List (CTL). This allows you to define a list of trusted root CAs and to specify the purposes for which you will trust those CAs (for example, secure e-mail). The CTLs are then deployed by using Active Directory Group Policy Objects (GPO). Although this method is convenient, it is proprietary to Microsoft. Only clients running Windows 2000 or later can use CTLs. This option only affects clients in the domain in which the CTL GPO is applied.
- Install the root CA certificate into the Active Directory Trusted Certification Authorities store (in the Configuration container). This creates an unconditional trust in the root CA (and all subordinate CAs) for all of your users and devices throughout the forest. However, you should be extremely careful before granting this type of trust. Use this method only for CAs that are under the control of your own organization.
- Deploy a trusted root CA certificate to a subset of users or computers by using Group Policy. This is similar to the previous option but allows you to be much more specific about who and what will receive the trusted roots (that is, the users or computers targeted by the GPO). This option only affects users in the domain in which the GPO is applied.
- Use the Microsoft Root Update service. This service is intended to allow commercial certificate providers to easily distribute new roots to large numbers of people. You should consider disabling this service on all your corporate systems if you are intending to regulate your trusted root CAs.
- Use Group Policy to disable third-party trusted roots. In contrast to the other items in this list, this is a means of restricting trust rather than increasing it. Every computer running Windows (and the users that use these computers) inherits a set of roots that are installed by default. (This is also common with other operating systems and Web browsers across a range of platforms.) You can use Group Policy to disable the automatic trust in these roots. You can then use one of the previously described mechanisms to selectively add back trusted roots that you need (with or without restrictions, depending on your security needs).

Note: There are certain root certificates that you cannot disable. This is because the operating system relies on them for things like driver signing policy. These required roots are not disabled by this Group Policy setting.

This solution disables the Update Root Certificates service on the CAs. You should consider disabling this service on the other computers in your organization. Also consider using Group Policy to disable the default third-party roots for all domain users. Chapter 7, “Implementing the PKI,” addresses these items in more detail.

The root CA certificate of the PKI in this solution is distributed to clients by importing it into Active Directory, as discussed in the next section.

Root CA Certificate Distribution

Root CA certificates are automatically distributed to Active Directory forest members. By importing the CA certificate into the Certification Authorities container, members (computers and users) of all domains in the forest will install this certificate into their local trusted root CA stores. This is the recommended method for all internal root CAs that need forest-wide trust scope.

You will typically also need to distribute roots with a more limited trust scope alongside your internal roots. For more information about this topic, see the “Extending the Certification Authority Infrastructure” section later in this chapter.

To distribute your root CA certificate to users and computers on other platforms or outside of the forest, you must install the certificate manually or use some other method of distributing the root certificate to them.

Defining Certification Authority Roles

Now that you have defined the trust model and selected the root CA strategy, you can plan the rest of the CA infrastructure. To do this you must define the different roles that the CAs will fulfill in your organization. You can configure the CAs as root CAs or subordinate CAs. Subordinate CAs can, in turn, be issuing CAs or intermediate CAs (in that they are intermediate trust steps between issuing CAs and root CAs).

Root CA

The root CA role is very important in any organization. It is a role that is explicitly trusted by all users and devices in your organization. Many security decisions (such as whether to allow someone to log on, trust an e-mail, or permit a \$10 million securities trade) ultimately rest on the security of this root and the private key that provides its identity. Because so many operations depend on the root, changing a root key and certificate can be a very complex and error-prone operation that can lead to intermittent loss of service for applications and users for an extended period.

For this reason, it is highly desirable to protect the root CA private key as much as possible. One of the best ways of doing this is to disconnect the CA from the network so that access to it is extremely limited (You should combine this protection measure with equivalent measures to restrict physical access to the server). A further enhancement to protecting a CA's keys is using a dedicated Hardware Security Module (HSM). These provide additional key security for offline CAs, as well as greatly improved security for online CAs.

The solution defined in this guide uses an offline root CA.

Important: You should consider using an HSM for your root CA to enhance the security of the CA keys. You can add HSMs after installing your CA, but it is much simpler and more secure to do this from the outset. If you install an HSM later, you should renew your CA with a new key even though many vendors allow you to import the existing key.

Intermediate and Issuing CAs

Taking the root CA offline makes it impractical to issue certificates from it on a day-to-day basis. To create CAs that you can use to issue certificates for day-to-day use, the root CA certifies subordinate CAs to issue certificates on its behalf. This allows a subordinate CA to inherit the trustworthiness of the root CA without exposing the root CA key to security threats.

This process can be taken further. Instead of issuing certificates directly, the subordinate CA instead certifies a further layer of CAs to issue to end entities (users and computers). This not only provides an additional security layer to the root CA key, but it allows you to partition risk between subordinate CA branches. For example, one intermediate CA can certify internal issuing CAs, while another intermediate CA certifies CAs that issue external certificates. This method has the following benefits:

- It helps to restrict the issue of CA compromise to a smaller part of the whole PKI hierarchy.
- It allows separate certificate policies to be implemented for whole branches of the CA hierarchy.
- It reduces the number of times that the root CA key is used, which reduces the opportunities for compromise of the key.

Although additional layers of intermediate CAs increase the overall security of the PKI, this does come at a cost of extra complexity, additional hardware and software, and increased management overhead (the latter typically being much greater than the hardware or software license costs). For many applications, the security requirements do not justify a three-tier hierarchy. This is particularly true if the CA keys are protected with HSMs.

The solution defined in this guide proposes a two-tier hierarchy. The solution design provides an acceptable balance between good security and affordability, while also providing flexibility for future certificate applications, (for example, see the details outlined in the “Defining Certificate Requirements” section earlier in the chapter).

Note: Discussion on government or regulatory requirements that state that you must use three-tier hierarchies is beyond the scope of this guide. These requirements obviously override other considerations.

Proposed CA Hierarchy

The following figure illustrates the proposed hierarchy. The current implementation comprises the root CA and one issuing CA. The issuing CA will primarily issue standard assurance (sometimes referred to as “technical”) certificates for computers or users and higher-value certificates for computers.

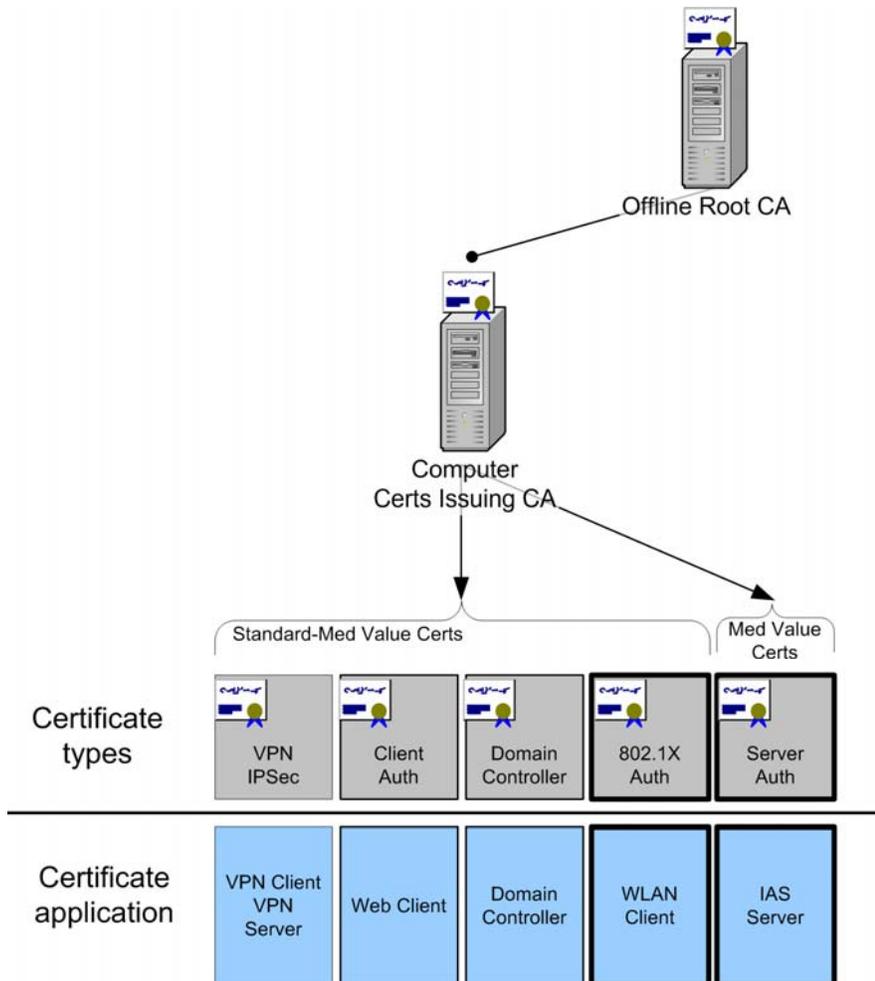


Figure 4.2
Certification Authority hierarchy

This design allows you to deploy a fully functional PKI capable of supporting secure wireless LAN authentication (802.1X) with a relatively small outlay on hardware, software, and management costs.

Note: You can extend this simple infrastructure design in several ways to accommodate different requirements. This topic is discussed in the “Extending the CA Infrastructure” section.

The issuing CA will be initially configured to issue the following types of certificates (these are shown in bold boxes in the previous diagram):

- Client Authentication–User
- Client Authentication–Computer
- IAS Server Authentication

The first two are standard certificates and are automatically issued based on domain credentials of the user or computer. Possession of these certificates indicates no stronger binding with the subject than possession of a valid domain user name and password. However, it is important to note that there are security and other technical advantages in using certificates in place of domain names and passwords.

The IAS Server Authentication is classified as a medium assurance certificate because the IAS servers perform a high security function. The issuance of this type of certificate will typically include extra checks into the validity of the request, and it will require certificate manager approval.

Note: In the later build chapter that describes creating this certificate type, the requirement for Certificate Manager approval is left disabled. This allows the IAS servers to automatically renew expiring certificates and avoids the service being disabled when the certificate expires. As long as you have processes in place to adequately vet and approve the certificate request, you should consider enabling the requirement for Certificate Manager approval.

Hardware and Software Requirements for CAs

This section discusses the hardware and software requirements for the CAs.

Root CA

The hardware requirements for the root CA are minimal. The computer specification needs to be no greater than the minimum requirement to simply run Windows Server 2003. The critical characteristics of the root CA hardware are long-term reliability and maintainability. The root CA typically resides on a computer with a long life span but one that is turned off most of the time. When the computer *is* turned on though, you want it to start reliably. To respond to possible hardware faults, you will want to be able to replace components easily, possibly years after the computer model has been discontinued.

With these consideration in mind, Microsoft recommends to:

- Choose a well known computer manufacturer with a good record for support and long-term hardware maintenance. Inquire about how long you will be able to obtain spare parts from the vendor.
- Use server rather than workstation or portable computer hardware because the former tends to be more standardized and changes less frequently.
- Consider maintaining a spare system that can take over the role of the root CA if the hardware fails and cannot be repaired in a short time.
- Keep a copy of the original installation media, drivers, and patches so that you can rebuild the system if it fails.
- Consider using an HSM for additional security.

The root CA does not require the additional capabilities of the Enterprise Edition of Windows Server 2003. For this reason, the solution in this guide uses the Standard Edition of Windows Server.

Issuing CA

Although there are performance requirements for the issuing CA, these are relatively low because the CA normally does very little work. Even when heavily loaded, performance measurements suggest that for an Enterprise CA, the interaction with Active Directory is usually the constraint (not the CA itself). Therefore, the hardware performance requirements are fairly modest. As with the root CA, reliability and maintainability are critical factors in your choice of hardware.

Certificate Services uses the same database technology as Active Directory, so many of the same performance guidelines apply. A good guideline for most organizations is to use the same hardware specification that you use for Active Directory domain controllers.

For more information about CA performance, see the “CA Capacity, Performance, and Scalability” section in the “Designing a Public Key Infrastructure” chapter referenced at the end of this chapter.

Chapter 7, “Implementing the Public Key Infrastructure,” gives a suggested hardware profile. In addition to the preceding guidelines for the root CA Microsoft recommends considering the following when specifying server hardware for the issuing CA:

- Redundant network interface cards (NICs) using NIC teaming.
- Two redundant array of independent disks (RAID 1) volumes are the recommended minimum so that you can store the CA logs on a separate physical storage unit. This adds performance benefits as well as resilience to hardware failure.
- Consider using three RAID 1 volumes (rather than two) to store the operating system, Certificate Services database, and Certificate Services logs, respectively, on separate physical volumes for better performance.
- High performance small computer system interface (SCSI) drives and controllers are preferred over integrated device electronics (IDE) equivalents for their better performance and resiliency characteristics. Aside from interaction with Active Directory, the performance of the disk subsystem is possibly the most significant factor in determining overall CA performance.
- Consider using an HSM both for additional security and increased performance of the signing operations during certificate issuing.

In contrast to the root CA, the issuing CA *does* require the additional capabilities of Windows Server 2003 Enterprise Edition to support editable certificate templates and user certificate auto-enrollment.

Using Multiple Issuing CAs for Service Resilience

This section discusses the technical reasons why you might want to install multiple issuing CAs. There also are security and policy reasons why you may want different issuing CAs to enroll different certificate types. These reasons are considered in a later section of this chapter.

A single issuing CA with very modest hardware is adequate for issuing the certificate types described earlier to tens of thousands of clients, so you are unlikely to need multiple issuing CAs purely for performance reasons. However, you should consider whether your availability requirements for the issuing CAs mean that you must deploy multiple CAs to enroll the same certificate types.

A CA does not have the same type of availability issues as many services. Clients do not need to contact the CA to use or verify a certificate. Clients only directly contact the CA when the CA is required to:

- Enroll a new certificate.
- Renew a certificate.
- Revoke a certificate.
- Publish a new CRL.
- Renew the CA certificate itself.

The availability requirements of each of these are detailed in the following table.

Table 4.8: CA Service Availability Requirements

CA service	Availability requirement
Enrollment Services — New Certificate	This may be significant factor given that it may prevent new users from accessing the network or other services requiring a certificate. You must assess whether the time required to recover the CA from backup is longer than your organization can wait for a new user to enroll a certificate. Most organizations judge that cost of waiting for the CA to be recovered is less than the cost of managing additional CAs. Otherwise you need multiple issuing CAs for the certificate types concerned.
Enrollment Services — Renew Certificate	If automatic renewal is used with the certificate type in question, this occurs by default six weeks prior to the previous certificate expiring. By contrast, recovery time from backup for a CA is usually measured in hours. Manually renewed certificates are left to the owner to renew. You may want to institute an automated warning system that will alert the owner when critical certificates are due for renewal. Otherwise, the availability criteria are the same as for enrolling a new certificate.
Revoke a Certificate	A certificate can normally only be revoked by the CA that issued it — a second CA would not help. If the revocation is extremely time-critical (that is, it needs to be done before the CA can be recovered), you can insert revocation entries into current CRLs as long as you have the serial number of the certificate to be revoked and the CA private key (restored to a different computer). Note that CRLs typically have a latency of one day or more. Unless the recovery time of the CA is longer than the interval to next CRL publication, you gain little by manually updating the CRL.

(continued)

Publishing a CRL	<p>A CRL is unique to a CA — a second CA would not make CRL publication more resilient; it would only minimize the impact of a CRL publishing failure (since less than 100 percent of the issued certificates would be dependent on the failed CRL).</p> <p>Access to current revocation status is essential for many certificate applications. This means that a CRL that has not expired must be available at the published CRL Distribution Points (CDPs). If this does not happen, revocation-sensitive certificate applications will fail.</p> <p>The CA recovery period should never be longer than the overlap between previous CRL expiration and new CRL issuance. In the rare cases where it is, a CRL can be resigned and have its validity period extended. This procedure is covered in the Operations Guide.</p>
Renewing the CA Certificate	<p>A second issuing CA cannot help with this task.</p> <p>This operation should never be left so late in the process that the recovery time of a CA becomes an issue. Even if it does, the CA certificate can be re-signed with the parent CA key to extend its validity period.</p>

Note: In the preceding table, CA recovery time and CA availability refer to anything that affects the CA's ability to deliver service to end users. This is not confined to server failure. In fact, network outage between sites is an example of a much more likely cause of service failure. You should consider all of the factors that may impact service delivery to users when deciding on your required level of service availability.

As long as you manage the backup and recovery of the CAs well, only the enrollment and some renewal requirements will affect your decision to use multiple CAs to provide service resilience. You must weigh the cost of these services not being available against the installation and management costs of providing additional CAs.

As well as improved availability, multiple issuing CAs also give better certificate issuing performance and halve the size of CRLs. None of these factors are significant in the solution for this guide. This solution handles resiliency issues by carefully managing the CAs, and by including adequate backup and recovery procedures. For these reasons, only a single issuing CA is required for this solution.

Protecting the CA Key with HSMs

One significant improvement that you can make to enhance the security of the basic solution presented in this guide is to use HSMs to protect the private keys of all CAs. Although these are often costly—they may easily cost more than the CA server—the added level of security that they bring to your environment is significant. Taking this measure allows you to restrict access to CA key operations to only authorized users. Sensitive operations (such as exporting and backing up the CA keys) are typically protected by multiple smart cards. This is more secure than only relying on software-based keys, which can be copied from the CA by any member of the local Administrators or Backup Operators groups.

Aside from the enormous security benefits of HSMs, they can also speed up CA operations by offloading work from the CPU to dedicated cryptographic acceleration processors.

CA Security

This section examines the security of the CAs, including operating system and physical security, security auditing and monitoring, and the use of roles to delegate management responsibilities over the CA.

Operating System Security

The CA is secured by using Windows security policies. The settings are based on the Certification Authority server role in the *Windows Server 2003 Security Guide*.

For more details on the settings used in this role, see Chapter 7, “Implementing the Public Key Infrastructure.”

The security settings for the root CA are applied directly using the security templates, whereas the issuing CA settings are applied using Group Policy.

Physical Security

The physical security of the CA servers is paramount. Unless you can control basic physical access to the servers, no amount of network or operating system security will be effective.

The root CA should reside in a location where access to the server is strictly controlled. Access to the CA is required only rarely (two to three times a year), and the server does not need to be powered on other than during these occasions. This means that you can store the server in a safe storage room that does not have the standard computer facilities of a server room in it. (For example, the storage room does not need networking, sophisticated server accommodation, or special power and temperature management.

The issuing CA should also reside in a location where physical access is strictly controlled. Physical security is important because there are many ways to compromise the security of a computer system if an attacker has physical access to it (over and above those attacks that are possible over a network). Because this server needs to be continuously online, you should store it in a location that has standard computer server room facilities (temperature control, power management, air filtering, and fire extinguishing capabilities).

If possible, choose locations for both servers that are as free from external risks that might damage them, such as, fire, flood, as they can be.

It is equally important to control physical access to, and ensure the physical safety of, backups, key material, and other configuration data. You should store this information at a different location than the servers themselves to allow for recovery of the CAs in the event that the whole site becomes unavailable (for example, following a natural disaster or a fire).

Security Management of the CAs

A certificate infrastructure is potentially a very high value asset. How high that value is will depend on what your organization uses certificates for — not just now but over the next five years or more. Because of this you should carry out the installation, configuration, and management of the CAs using more stringent security and verification measures than those you might use for installing other IT infrastructure. These measures should be at least equivalent to those designed for a domain controller. In some cases you may find that you need a higher level of security than this.

The trust that you place in a CA depends on you having a high level of confidence that it has been set up and managed securely. If you cannot guarantee that the CA's private key has not been surreptitiously copied, you can never be really certain that a certificate supposedly issued by that CA is not a forgery.

This assurance or confidence level cannot be easily increased retroactively; you must build this special status into all interactions with the CA from the beginning. For example, your organization's confidence in the fact that the CA's private key has not been compromised will be a lot higher if you have some audit trail or other evidence that all accesses to the CA have been legitimate. For example, all administrative operations on the CAs have been witnessed by a person other than the administrator or recorded on video. In the case of an offline CA, the fact that it has never been connected to a network greatly reduces the possibility that it can have been compromised.

The need for this high level of assurance may be particularly relevant if your organization ever enters a legal dispute over the validity of a certificate that it has issued. In such cases, if you have evidence that the CAs have not been compromised, you have a much greater chance of a favorable outcome to the dispute. A full discussion of this subject is beyond the scope of this guidance, and you should consult your auditors and legal advisors to explore this subject further.

Some examples of the kinds of steps that you can take to significantly increase the assurance level of your CAs are:

- Ensure the physical security of the CA so that unauthorized individuals cannot access the CA hardware or backup media.
- Perform all installation and configuration steps with a witness present—record the major installation steps and have the witness countersign to verify that these have been carried out successfully. (An alternative to this would be to videotape the installation and entrust a copy of the video to a trusted party).
- Perform all certificate issuing and revocation operations on the root CA under similar conditions. Ideally, all access to the root CA should be witnessed.
- Ensure that all individuals who have administrative access to the CAs have individual accounts uniquely traceable to them. Audit all operations on the CA.
- Consider enabling role separation on the CA. (This is discussed in more detail later in the chapter.)

These types of measures are particularly important for the root CA server. The issuing CA can have a much lower assurance level depending on the types of certificates that it needs to issue. For example, if the CA is not issuing high value certificates (only standard certificates such as computer and user network authentication), you will not need any greater security for this CA than you use for a domain controller.

As long as the root CA has a high assurance level, you have the flexibility of adding a higher-assurance issuing CA to issue higher value certificates later. You can maintain high assurance CAs alongside the existing, standard CA. However, if the root CA is installed and configured in a relatively low-security environment and you later want to issue high value certificates, you will probably need to reinstall it or create a new root CA.

Security Monitoring and Auditing

Operating system and Certificate Services auditing is used on all of the CAs. To be effective, you must monitor the auditing and any suspicious items. See Chapter 11, "Managing the Public Key Infrastructure," for discussion on the significance of Certificate Services audit event entries.

Management Roles

Certificate Services gives you a great deal of control over the delegation of administrative roles. This solution uses this capability to provide you with a lot of flexibility in how you administer the PKI. Each of the core administrative roles defined by Certificate Services has been implemented by using a domain security group or, for offline CAs, a local security group. In addition, two more roles and security groups have been defined in this solution to help with delegation of administrative duties over the PKI components of Active Directory.

It is important to understand that there is not necessarily a one-to-one mapping between these roles and different IT personnel within your organization. Most organizations will find that the same person fills many of the available roles. You can easily implement this by adding that person to any or all of the security groups listed in the following table. Conversely, if your organization has a more complex separation of administrative responsibilities, the capability is already there to implement this.

The implemented roles and their mappings to security groups (where implemented) is shown in the following table.

Table 4.9: Core Certificate Services Roles

Role Name	Security group	Scope	Description
Enterprise PKI Administrator	Enterprise PKI Admins	Active Directory Forest	Responsible for overall PKI — defines certificate types, application policies, trust paths, and so on for the enterprise.
Enterprise PKI Publisher	Enterprise PKI Publishers	Active Directory Forest	Responsible for publishing trusted root certificates, sub-CA certificates, and CRLs to the directory.
CA Administrator	CA Admins	CA	Responsible for CA configuration. Often will be the same individuals as those in the Enterprise PKI Administrator and Administrator roles. There may be multiple CA administrators in charge of different CAs if the certificate usage dictates this.
Administrator	Local Administrators	CA	Administers the CA operating system and server. Responsible for installing the CA and for renewing the CA certificate. Typically shared with CA Administrator role.
CA Auditor	CA Auditor	CA	Manages the audit event log and policy of auditable events from the CAs.
Certificate Manager	Certificate Manager	CA	Approves certificate requests that require manual approval and revokes certificates. There may be multiple Certificate Managers in charge of approvals on different CAs if the certificate usage dictates this.

(continued)

Registration Authority or Enrollment Officer	Not defined	Certificate Profile	This is an extension of the role of Certificate Manager. Responsible for approving and signing certificate requests following out-of-band ID verification. Can be a person or an IT process or device (such as a fingerprint scanner and database). You can specify different Registration Authorities for different certificate profiles (templates) and can span multiple CAs.
Key Recovery Agent	Not defined	CA	Holds key to decrypt archived private keys in CA database.
CA Backup Operator	CA Backup Operator	CA	Responsible for backup and recovery of CA servers and secure storage of backup media.

These security groups are implemented as domain universal groups and applied to the issuing CA and directory. For the root CA, equivalent groups are implemented as local groups (although there is no equivalent for the Enterprise PKI Admins and Enterprise PKI Publishers for an offline CA). The solution assumes that the same security groups will be used for all CAs within the enterprise. If this is not valid for your organization, you should implement separate groups for each CA for all of the roles with CA scope. (Obviously, these must be renamed appropriately, for example CA Admins — Issuing CA 1.)

Since Registration Authorities (or Enrollment Officers) and key archival and recovery have not been implemented for this solution, these roles do not have defined security groups.

It is possible to enforce separation between the CA roles for a CA. When this is enabled, any user who is a member of more than one role group is denied access to the privileges of all role groups. Role separation is not implemented in this solution.

Active Directory Integration

You can install CAs in one of two modes — Enterprise mode (or Active Directory Integrated) or Stand-alone mode. The key differences between the two modes are that Enterprise CAs: rely on Active Directory to store configuration information; can use Active Directory as a registration authority; and can automatically publish issued certificates to the directory. Stand-alone CAs can publish certificates and CRLs to the directory, but do not rely on the presence of Active Directory.

See the “More Information” section at the end of the chapter for resources that provide a more detailed explanation of this.

Since it is offline, you can only configure the root CA as a stand-alone server. The same is true of offline intermediate CAs if you plan to deploy them in your environment.

The issuing CA will be configured as an Enterprise CA for the following reasons:

- Certificate auto-enrollment and auto-approval is required for the certificate types used in the solution.
- Certificate templates are required in the solution—they provide considerable benefits by easing the management of multiple certificate types (often across multiple CAs).
- IAS requires Active Directory to perform trusted certificate mapping to authenticate the wireless clients. The CA must be registered in the NTAAuth store to permit this.
- Automatic publishing of certificates to corresponding user or computer objects is possible (although not required in this solution).
- A trusted source to obtain subject name information for use in the certificate requests and issued certificates is required by the CA—Active Directory can provide this from the user and computer attributes stored in the directory.
- Smart card logon certificates may be required in the future—these are much easier to implement using Enterprise CAs.

Note: While these capabilities are provided by default on an Enterprise CA, you can also provide some of them from by a correctly configured stand-alone CA. This is described in more detail in the Certificate Services product documentation. See the reference at the end of this chapter.

Installing CAs in Your Domains

If your organization has a multidomain forest (or even multiple forests), you must decide on the domain or domains in which you will install your CAs. Your decision may be affected by many factors, such as the need to delegate control to different domain managers, or national or regional legislation affecting the provision of certificates to different parts of your organization.

The most common approaches are to install the CAs into the forest root domain or into a dedicated domain for management purposes. You should install them into a domain that will remain stable over a long time. (You cannot change a CA's name, domain membership, and DNS domain name after installation.) You should also avoid installing the CAs into domains where you cannot guarantee the security or integrity of the computer account. Although it can make centralized management easier, you do not need to install your CAs into the same domain.

In this solution, the CA server accounts (Enterprise Issuing CAs only) are installed into the forest root domain, or if it is a single domain forest, into this domain.

Mapping Your Certificate Practices Statements to CAs

If you are planning to publish your CPS, you must determine the scope of the CPS. You can create a CPS for a whole CA hierarchy or part hierarchy, or you can have a CPS per CA.

The latter option gives you the most flexibility, but it also increases the overhead of managing multiple CPSs. Standard practice is to create a separate CPS for each CA or group of CAs that have a common certificate policy, subject types, and security levels. Where these differ significantly between CAs, you may need to use multiple CPSs. Obviously, if you have many identical CAs deployed for resilience or performance reasons, these should have identical CPSs.

As discussed previously, it is quite legitimate to create a CPS, but not publish it. For example, you may want to avoid publishing your CPS externally if you feel that it contains operational and security information of an internal nature. You may also decide to publish an abbreviated version of your CPS that lists the important operating policy of the CA without disclosing any internal operational detail.

If you decide that you will publish your CPS and want to advertise its location in the CA certificate, you must obtain an object identifier (OID) for your certificate policy from the official OID namespace assigned to your organization by the International Standards Organization (ISO). Your certificate policy is unique to your organization, so it needs a globally unique OID to identify it.

This CP OID is encoded into each of your organization's certificates as a certificate extension. A “*certificate extension*” is a type of certificate data field. A URL is included as part of this extension that points to the CPS for the CA that issued a given certificate.

It is common to also include a user notice as text that gives some indication of the purpose or origin of the certificate (although this is limited to 200 characters, so it is obviously not an alternative to a separate CPS document).

For precise details on how certificate policy OIDs and CPS URLs are encoded into a certificate, see RFC 3280, *Internet X.509 Public Key Infrastructure Certificate and CRL Profile* in the “More Information” section at the end of this chapter.

Once you have obtained your CP OID and decided on the URL where the CPS will be published, you can include this in your CA certificates. Chapter 7, “Implementing the Public Key Infrastructure,” documents this procedure.

Supporting IT Infrastructure

The PKI in this solution relies on other infrastructure services to operate correctly. The following diagram illustrates the principal ones — Active Directory and IIS — and how they interact with the CAs and certificate clients.

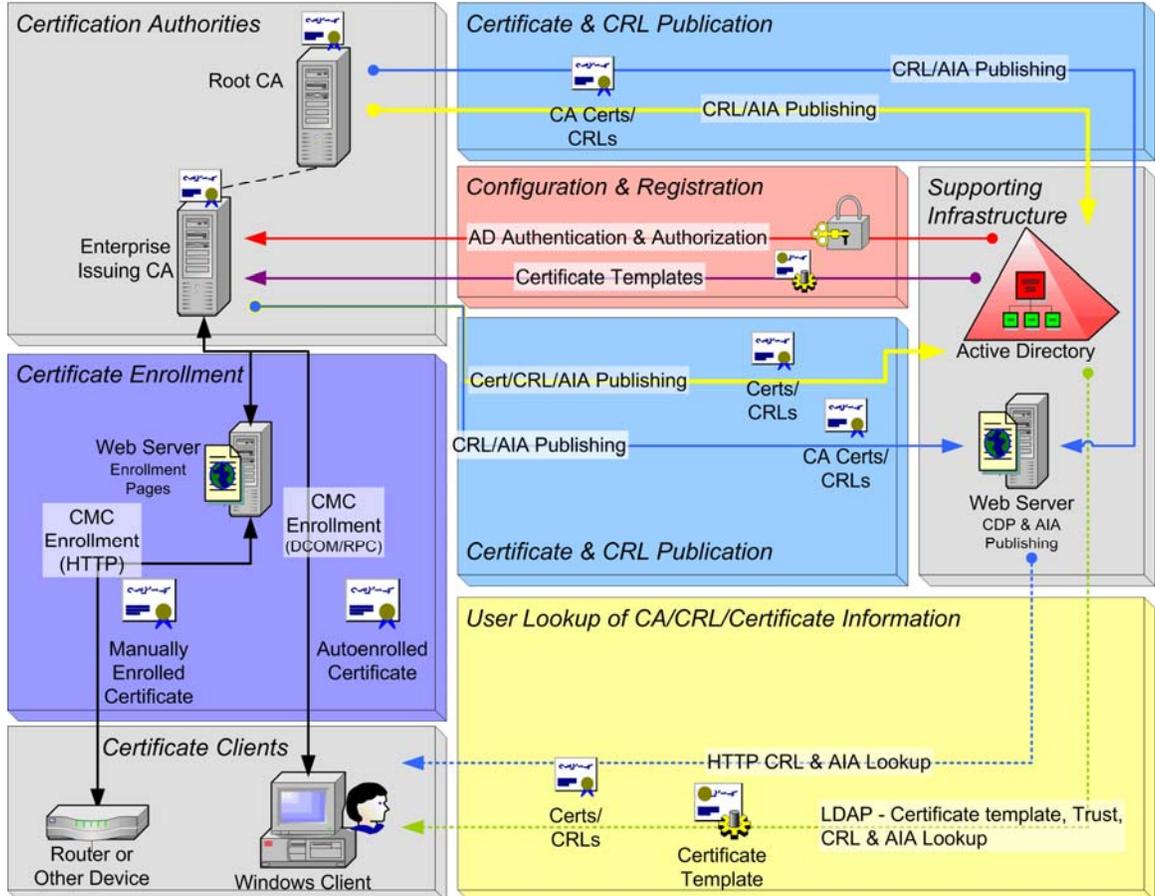


Figure 4.3
CA interaction with IT infrastructure

Monitoring, alerting, and management infrastructure is also pivotal to the successful operation of Certificate Services, although these components are not illustrated in the diagram. Chapter 11, “Managing the Public Key Infrastructure,” describes this infrastructure in more detail. The following sections describe the services that Active Directory and IIS provide to the PKI.

Active Directory

As discussed in the earlier section, “Active Directory Integration,” Active Directory provides a number of different services essential to the PKI. These include certificate publishing, certificate registration, certificate account mapping, and storage of trust and configuration information.

All of these services are provided automatically for Enterprise CAs. Online stand-alone CAs can also use some of these services. However, CAs that are offline cannot directly interact with Active Directory to store and retrieve information.

In this solution, the CA certificate of the offline root CA is published to Active Directory trusted root store. This results in the root CA certificate being automatically distributed to the trusted root store of all Active Directory clients within the forest.

The root CA could also use the Active Directory for the following publishing services (although these are not used in this solution):

- CRL publishing —domain clients (throughout the forest) can retrieve CRLs from a local domain controller that have been published to Active Directory.
- Cross certificate distribution to domain clients—Cross certificates published to Active Directory are automatically distributed to the local certificate store of each Active Directory client in the forest.

The issuing CA uses Active Directory for all of the services described in the “Active Directory Integration” section.

Using Active Directory for Non-Active Directory Clients

There are a few considerations for supporting PKI clients who are either members of another, untrusted Active Directory forest or are not members of any Active Directory forest.

If you want to allow these types of external clients to retrieve CA certificates and CRLs using Lightweight Directory Access Protocol (LDAP), you must consider the following items:

- External clients require that you configure an explicit LDAP host name for CDP and AIA paths. For more information, see the “Configuring CDP and AIA Paths” section later in this chapter.
- External clients will not perform anonymous LDAP queries on Active Directory by default. You must change the **dsHeuristics** value of the forest and grant explicit access permissions to the Anonymous account. For further information about this topic, see the “More Information” section at the end of the chapter.

Warning: This allows anonymous LDAP on all domain controllers in the forest (although only items with explicit permissions for the Anonymous account will be accessible by unauthenticated clients). Carefully consider the implications of allowing unauthenticated access to your directory before doing this.

- External clients will not inherit the trusted root information configured in the directory—you must configure this information by some other means.

Considerations for Internet Clients

There are also some important considerations for CDP and AIA configuration for clients outside your organization (such as Internet clients). This is covered in the “Configuring CDP and AIA Paths” section later in this chapter.

If you need to provide certificate lookup for Internet clients to support, for example, e-mail certificate lookup, you may need to create a separate LDAP directory for Internet clients. Because of the security implications, Microsoft strongly recommends to not use the method described in the previous section for enabling anonymous LDAP to your internal Active Directory forest. Instead, create a separate Active Directory forest, and replicate information from your internal forest by using the LDIFDE tool, a metadirectory product (such as Microsoft Metadirectory Services), or another directory synchronization product.

Internet Information Services

IIS provides two services for the PKI in this design:

- It publishes CA information such as CRLs, CA certificates, and potentially CPS documents.
- It provides the ability to enroll certificates using a Web interface—especially useful for non–Windows clients—although this capability is not used in this solution.

Using IIS to Publish CA Information

In this solution, CDP and AIA information for both the root and issuing CAs is published to a Web server. Using HTTP publishing allows the broadest range of clients to retrieve the required information.

The installation of IIS on the issuing CA for this purpose is common but may not always be the best approach. If you can use another IIS server to publish the CRL and CA information, use it instead. Try to limit the ways in which users can access the CAs because every additional protocol and service on the CA gives an attacker another possible entry point to the server. Using IIS on the CA also hinders the CA being shut down for maintenance, since this may be the only valid CRL location for many clients.

In the Build Guide, the issuing CA is used to host IIS for CRL and CA publishing. This was done to simplify the build process and reduce the requirement for extra hardware. However, Microsoft recommends locating them separately if you can.

Using IIS Enrollment Pages to Enroll Certificates

IIS enrollment pages are useful in a number of scenarios: for nondomain client enrollment, non–Windows client enrollment, or for browser clients other than Microsoft Internet Explorer.

However, the Web enrollment pages are not required on a CA. They are installed on the issuing CA in this solution by default, although you can install them on a separate Web server instead (the server must be running IIS 5.0 or later to support ASP pages). The enrollment pages make some enrollment tasks easier, but you do not need to install them at all if you do not need them. For more information about installing and using the Web enrollment pages, see the “More Information” section at the end of this chapter.

Configuring CDP and AIA Paths

Clients need access to an up-to-date CRL to determine whether a certificate has been revoked. Clients also may need to retrieve CA certificates to verify that an end-entity certificate chains to a trusted root. Each CA needs to encode into its certificates one or more URLs that provide locations clients can use to obtain information about the certificates.

What you configure the CDP and AIA for each of your CAs depends on the types of the clients that will use your certificates. Are the certificates only intended for use by users and computers who are members of your Active Directory forest, for example? Or do external users or devices also need to use them? Defining certificate clients is covered earlier in this chapter.

Root CA

The root CA is configured for this solution as follows:

- The primary (first-listed) paths for the CDP and AIA are HTTP URLs. The root CA CRL is usually very small (1–2 KB), and the publication interval very long (six months), so publishing it in a single location will not be a significant constraint to clients retrieving the CRL.
- The secondary paths are configured as LDAP URLs to create a backup to the HTTP locations. No LDAP host name is used, so clients in the same forest will retrieve the information from their local domain controllers. Clients outside the forest cannot access this location.

This configuration allows clients outside the Active Directory forest to use certificates from this CA and subordinate CAs because they will default to using the HTTP paths.

Issuing CA

The issuing CA in this solution is optimized for use by internal Active Directory clients. The CA is configured as follows:

- The primary paths for the CDP and AIA URLs are LDAP directory paths.
- No LDAP host name is specified in the CDP and AIA URLs; this lets the client use its default LDAP server. In the case of Active Directory clients, the default LDAP server is their local domain controllers. However, other LDAP clients may fail when querying these LDAP paths.

This configuration is optimal for clients in the same forest as the CAs. Base CRLs are published weekly, and delta CRLs are published daily. Since the default location for both of these is Active Directory, clients retrieve these from the nearest domain controller. This provides resilience and optimizes network traffic.

Setting CDP and AIA to Support External Clients

The previous arrangement is not optimal for clients who are either members of another Active Directory forest or not members of any Active Directory forest (for example, a router). Because these foreign clients cannot access the LDAP CDPs and Authority Information Accesses (AIAs), external users can experience significant delays in trying to check revocation and AIA information. These delays may cause applications to fail for these external users.

If there is the possibility that the certificates will be used by clients outside your Active Directory forest you must configure the CDP and AIA values to use HTTP URLs as the primary paths instead of LDAP URLs.

To include LDAP URLs that are usable by external clients, you must do the following:

- Configure an explicit LDAP host name for CDP and AIA paths – you cannot use the default null path (LDAP:///). To change a CDP or AIA path for a CA requires reissuing (renewing) the CA certificate.
- Enable anonymous LDAP access as described earlier in this chapter.

Considerations for Internet Clients

If you are planning to distribute certificates outside of your organization for use on the Internet, there are a few additional considerations.

Certificates with internal LDAP URLs could provide information about the internal Active Directory and CA structure and names. To avoid this, you should:

- Only use HTTP URLs for the CDP and AIA values of the root CA and for any subordinate CAs in the chain.
- Issue certificates for external use from a separate CA. This CA will use only HTTP CDP and AIA URIs.

In both cases, you should provide a secondary HTTP location for CRL retrieval so that clients can fall back to this option if the primary location is unavailable.

For further discussion on using CRLs and CDPs, see the references included in the “More Information” section at the end of the chapter.

Extending Your CA Infrastructure

The “Defining Certificate Security Requirements” section discussed categorization of certificates by security level and also by subject type. The principal reason for separating different subject types is that it is very likely that different certificate policies and operational practices (as documented in the CPSs) apply to these subject types.

Typically, a CPS is applied per CA. It is possible to accommodate different policies governing different subject types in a single CA, but this can make the CPS complex and often difficult to implement correctly. The strategy for extending this PKI to issue certificates covered by different policy and security requirements is to create additional issuing CAs for the major subject types. This is illustrated in the following figure.

Note: This figure only indicates how you might extend the earlier CA hierarchy. It may be that your organization requires either something much more complex or simpler for your future needs. Base the design of additional CAs and certificate capability on your security requirements—there is no absolute right or wrong design for a PKI. When looking at extending the PKI to cover your other certificate needs, you should follow a similar approach to the one outlined here for the simple PKI design.

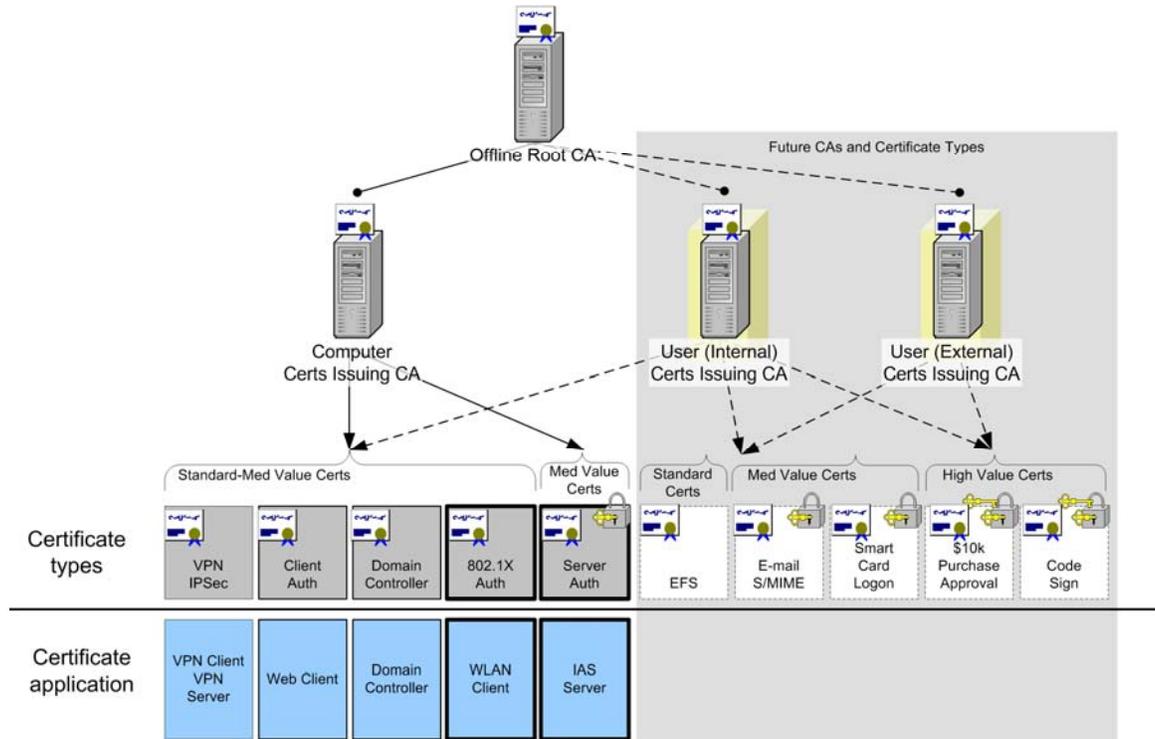


Figure 4.4
Extending the CA hierarchy

This diagram displays how you might extend the simple CA hierarchy presented earlier in this chapter to cope with a broader variety of certificate requirements. The new CAs and certificate capabilities are shown against the gray background. The diagram also illustrates the use of high value certificates (key and lock symbol) and how, as the user certificate CAs (Internal and External) are brought online, they will take over the role of issuing standard user certificates from the first CA.

This strategy for extending the CA infrastructure contains some assumptions:

- The CA infrastructure management will be centralized—that is, there is no requirement to delegate the control of CAs by organizational or geographical division.
- Common certificate standards are used throughout the organization—that is, a certificate of a given type has commonly accepted and agreed usages and policies throughout the organization.
- No interoperability with an existing PKI is required.
- You require different security levels and policies for the different certificate types shown (and any others that might be required).

If these assumptions are invalid for your organization, you may require a different structure than this one. For a detailed discussion of different options and approaches to extending your CA infrastructure, see the “Designing a Public Key Infrastructure” reference at the end of this chapter.

Qualified Subordination

You can use the certificate definitions that you created to define certificate templates and issue certificates without any further customization. However, in extending your PKI you might want to limit the scope of the certificates that your CAs can issue by constraining the delegation of your issuing CA certificates.

Qualified subordination, which was discussed in earlier in this chapter, can be used to control the scope and purpose of trusts within your organization by creating cross-certificates between your CA infrastructure and those of external organizations. You can use this same technique to limit the certificate types and certain attributes of certificates within your own CA hierarchy. Further discussion of this topic is outside the scope of this chapter. For more information about this topic, see the reference to the paper, *Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003*, at the end of this chapter.

For the PKI in this solution, qualified subordination within the CA hierarchy is not required.

Configuring Certificate Profiles

This section discusses how certificates are configured to address the requirements defined earlier in the chapter.

Defining Certificate Parameters

For each certificate type that you need, you should document the certificate profile for that type. You can then configure the profile parameters into certificate templates, which control the certificate types that will be issued by your CA.

Note: Stand-alone CAs do not use certificate templates. You must create the request using a tool such as Certreq.exe, by using the form on the Web enrollment pages, or constructing the request programmatically. If you use a standalone CA you should still define certificate profiles for each certificate type and use the profiles when building the certificate requests to a stand-alone CA.

A certificate profile definition includes all of the following:

- Template name and display name (You should define a naming standard for these.)
- Certificate key length
- Certificate validity period
- Optional certificate extensions
- Enrollment and renewal policies
- Policies related to validity periods
- Policies related to application usage
- Policies related to key usage
- Policies related to key archiving
- Certificate authorization
- Subject name creation
- Certificate enrollment agents
- Key creation
- Key and CSP types

Key length, validity period, key creation options, and the enrollment and authorization policies are all determined by the required certificate security level and application requirements identified earlier in the chapter.

Defining Certificate and Key Lifetimes

A number of factors affect certificate lifetimes, such as the type of certificate, your organization's security requirements, the standard practices in your industry, and government regulations. In general, longer keys support longer certificate lifetimes and key lifetimes.

There are constraints on key length and key type:

- **Compatibility**—some certificate applications may not support keys longer than 2048 bits. The key type may also impact compatibility; generally, RSA keys have the best compatibility but other key types may be required for some applications. You must consider both key length and key type compatibility for all CAs in the chain because applications will be required to process all certificates in the chain.
- **Performance**—signing and encryption operations require greater processing power for larger keys than for smaller ones. For this reason, keys larger than 2048 bit should generally not be used on issuing CAs with high rates of certificate issuance.
- **Storage**—larger keys create larger certificates that require more storage in the certificate database. If the certificates are published to Active Directory, its storage requirements will also increase. The size and time of backups will increase proportionally.

When you choose certificate and key lifetimes, you must consider how vulnerable your keys are to compromise and the potential consequences of such a compromise. The following factors impact the lifetimes that you choose for certificates and keys:

- **Private key length**—longer keys are more difficult to break so they justify longer key lifetimes.
- **CA security**—the more secure the CA and its private key, the longer the safe lifetime.
- **Use of specialized cryptographic hardware**—smart cards and HSMs make the private key more secure and justify a longer lifetime.
- **Trust in the certificate subjects**—you may allow longer certificate lifetimes for your employees and internal computers than for external users and computers.
- **The number of certificates signed using a CA key**—the more that the key is accessed and the more widely distributed the CA public key becomes, the more likely that the key will have been attacked and possibly compromised.

This key lifetimes and renewal periods for the CAs and end entities used in this solution appear in the following table.

Table 4.10: Certificate and Key Lifetimes

Certificate subject	Key length	Key lifetime	Renewal interval
Root CA	4096 bits	16 years	Eight years
Issuing CA	2048 bits	Eight years	Four years
End entity	1024–2048 bits	Six months to two years	90 percent of validity period

In terms of key length, 1024 bit keys are currently considered to be beyond practical cryptanalysis. These should have a safe key lifetime well in excess of the proposed two-year end-entity value. Keys of 512 bit length are generally no longer considered safe to use except perhaps for applications that have very low security requirements. For this reason, 512 bit keys are not used in this solution.

The issuing CA key strength is a compromise between security and performance requirements. A key of this size currently has a key lifetime of well beyond the four year renewal period.

The root CA has no real performance constraints, so you can set the key strength to the maximum of 16 kilobits. For compatibility reasons it is set to a much lower level in this solution. Even so, 4096 bit keys have a safe key lifetime well beyond the renewal period of eight years.

RSA keys are used by all CAs, although the key type for end-entities will be determined by their application requirements.

Although it is possible to renew a certificate with the same key, this is not recommended under normal circumstances. In this solution, a new key pair is generated at each certificate renewal.

Certificates issued by CAs cannot have a validity period that exceeds the remaining validity periods of the issuing CA, and all superior CAs up to the root. For example, if the CA certificate is going to expire in six months, you can only issue certificates with a maximum lifetime of six months. This solution therefore prescribes that the CA certificates are renewed after half of their certificate lifetimes. All certificates issued by a CA should have a validity period that is no longer than half of the issuing CA certificate lifetime.

This gives nested maximum validity periods of four, eight, and 16 years for end entities, issuing CAs, and root CAs, respectively. In this solution, end entity certificates are kept to a maximum validity period of two years. This allows an additional intermediate CA layer to be introduced without having too much impact on the certificate lifetime hierarchy.

Mapping Security Requirements onto Certificate Parameters

The following table lists the certificate types identified previously in this chapter and how the security category for each type maps to certificate profile parameters.

Table 4.11: Certificate Parameters

Certificate type	Issuance policy	Approval method	Key	Validity period	Key storage	Key export	CSPs
Client Authentication — User	Low	Automatic (domain auth)	1024	One year	Software	No	Named
Client Authentication — Computer	Low	Automatic (domain auth)	1024	One year	Software	No	Named
IAS Server Authentication	Medium	Manual (Cert Manager)	1024	One year	Software	No	Named

Notes:

The “Low” policy listed in the **Issuance Policy** column refers to the predefined “Low Assurance” certificate policy in Certificate Services. It is equivalent to the standard assurance or security level referred to earlier in this chapter.

The “Named” value in the **CSPs** (cryptographic service providers) column is meant to indicate that the CSPs allowed by that certificate type should be specified in the template and not left for the client to decide. The client computer and server certificates have specific CSP requirements.

Mapping Certificate Requirements onto Certificate Template Parameters

Applications will often expect certificates to be configured in a precise way. They may require that the subject name is formatted in a certain way, that specific application policy OIDs are included, or that the certificate has been issued with a specific issuance policy. If nothing else, the application will at least require that the key usage has been correctly defined. You must obtain all of these parameters from the application owner (or vendor) in order to define your certificate profile.

The application requirements for the list of required certificates in this solution are shown in the following tables. These tables illustrate both certificate properties and CA parameters (as configured in certificate templates). Not all possible properties are listed.

Note: Each of the certificate types detailed below is closely based on one of the built-in template types. Rather than edit the original templates, make copies of the built-in templates and edit the copies to arrive at the required settings. This allows you to easily revert to the built-in templates if needed, knowing that they have not been modified.

Table 4.12: Client Authentication — User

Certificate parameter	Value required
Certificate Template Name	Client Authentication - User
Active Directory Publication	No
Key Usage	Digital Signature
Key Archival	No
Minimum Key Size	1024
Subject Name	Common Name
Subject Alternative Name	User Principal Name
Application Policies/Extended Key Usage	Client Authentication
Cryptographic Service Providers (CSPs)	Microsoft Base Cryptographic Provider v1.0 Microsoft Enhanced Cryptographic Provider v1.0
Derived from which template	Authenticated Session

Table 4.13: Client Authentication — Computer

Certificate parameter	Value required
Certificate Template Name	Client Authentication - Computer
Active Directory Publication	No
Key Usage	Digital Signature Key Encipherment
Key Archival	No
Minimum Key Size	1024
Subject Name	Common Name
Subject Alternative Name	DNS name
Application Policies/Extended Key Usage	Client Authentication
Cryptographic Service Providers (CSPs)	Microsoft RSA SChannel Cryptographic Provider
Derived from which template	Workstation Authentication

Table 4.14: 802.1X Server Authentication

Certificate parameter	Value required
Certificate Template Name	802.1X Server Authentication
Active Directory Publication	No
Key Usage	Digital Signature Key Encipherment
Key Archival	No
Minimum Key Size	1024
Subject Name	Common Name
Subject Alternative Name	DNS name
Application Policies/Extended Key Usage	Server Authentication
Cryptographic Service Providers (CSPs)	Microsoft RSA SChannel Cryptographic Provider
Derived from which template	RAS and IAS Server

Creating Certificate Templates

Active Directory in Windows Server 2003 contains a set of predefined certificate templates for many common functions. When you install an Enterprise CA, it is configured by default to issue several of these built-in certificate types. You will find descriptions of all of the built-in templates in the Windows Server 2003 Enterprise Edition product documentation. (See the “More Information” section at the end of the chapter for a precise reference.)

In the solution presented in this guide, most of these default templates are removed from the CA; that is, they are removed from the templates folder in the Certification Authority management console — you should not delete the template definitions from the directory.

You can opt to use the predefined templates if these match your needs; most of the Windows certificate-based applications (such as Encrypting File System, VPN authentication, and others) are covered by these templates. If you need to issue other types of certificates, it is usually better to create templates that specifically align with your requirements. If you use the predefined templates without understanding their capabilities you risk enabling functionality that you did not intend. For example, the Computer certificate that is intended for simple client authentication can also be used as a Web server certificate.

To create new templates, you should find a predefined template similar to your certificate profile requirements, and then create a duplicate of the template to base your new templates on. Configuring the templates is a simple process of selecting the attributes to match the certificate profiles that have been defined in this section.

Note: You cannot create a new template from scratch; you make a copy from an existing template and edit that as needed. Computer templates must always be derived from computer templates, and user templates from user templates — the two are not interchangeable.

As you create and amend templates, keep a detailed record of the template parameters as part of your configuration management system.

Creating a Certificate Management Plan

After you have configured certificates for your organization, you must create a plan for managing certificates throughout their lifetimes. Creating a certificate management plan involves making decisions about the following:

- How requests for new certificates and certificate renewal are processed
- How to map certificates to user accounts
- How you manage and distribute CRLs
- The strategies that you use to recover encrypted data

Selecting Enrollment and Renewal Methods

You can perform certificate enrollment using several different methods (you also can renew certificates using all of these methods):

- Windows auto-enrollment.
- Online enrollment using the Certificate Enrollment Wizard (usually started from the Certificates Management Console).
- Online enrollment using the CryptoAPI or CAPICOM interfaces from applications or scripts.
- Using the Certreq.exe tool to create and submit requests and retrieve issued certificates.

Note: The previous four methods are actually just different interfaces to the same online enrollment interface.

- Web page enrollment.
- Manual offline enrollment (This involves generating the request as a file using one of the previously defined methods, taking it to the CA, submitting the request by using the Certification Authority management console, and retrieving it by using the management console.)

All of these methods are valid and appropriate for use in some circumstances. This solution uses the following enrollment methods:

- Windows Auto-enrollment. Where possible, this method is preferred for lowering the certificate management overhead. You can use auto-enrollment even where a certificate requires manual approval (but not where an enrollment agent signature is required). Even though the certificate will not be issued immediately, a request will be submitted to the CA, and the enrollment will complete after the request is approved.
- Manual offline enrollment. This method is used for all certificate enrollment and renewal at the root CA.

Neither of these methods is adequate for some more complex scenarios, for example, where a certificate request needs to be signed by a third party before submission to the CA. RPC-based online enrollment is also not supported on most non-Windows platforms (for example, routers). Auto-enrollment is also not possible in any case where you need to define the subject name or alternate subject name in the certificate request (rather than allow Active Directory to generate it).

To accommodate more advanced requirements such as these, consider using one of the following methods:

- Create a CAPICOM script to run on the client stand-alone computer, for example as part of a login script.
- Use certreq.exe in a command or batch file to generate and submit requests, and to retrieve and install the issued certificate.
- Create a custom Web (ASP or Microsoft ASP.NET) page using CAPICOM to build and submit the request. With this latter technique, it is possible to provide enrollment services for a wide range of clients and include sophisticated multistage processes (such as when you require multiple signatures to approve a request).

Mapping Certificates to Identities

The mapping between certificates and the subjects named in the certificates is a large topic of discussion that is beyond the scope of this chapter. However, there are two important aspects to consider on this subject:

- How is the identity of the certificate subject confirmed before the certificate is issued?
- How is the identity of the certificate subject discovered from the information supplied in the directory?

The first of these questions concerns how the certificate registration process is carried out. This aspect is covered in the next section, “Creating Certificate Policies.” The second question covers how certificate users (applications and services) correctly map the identity of the certificate subject to another identity that they can use. Examples of the latter are:

- How does a domain controller identify a user from his or her smart card certificate to log the user on to the domain and build an access token?
- How does an e-mail user discover the certificate of a person to whom they want to send a secure e-mail?

The majority of certificates are automatically mapped to Active Directory security principals (users and computers) as part of the enrollment process. Active Directory defines the Subject Name and Subject Alternative Name of the certificate to create an implicit mapping between the certificate and the security principal named in the certificate. The Subject Alternative Name will contain the UPN or e-mail name for a user and the Service Principal Name (SPN), or the DNS host name for a computer or computer process. The UPN and SPN values are unique within an Active Directory forest. The e-mail and DNS names should be globally unique although Active Directory does not enforce this. Other Windows services such as IIS and IAS can perform certificate mapping to user or computer identities as well.

Note: Certificate mapping is unrelated to certificate publishing. Certificate mapping signifies that there is some attribute of the certificate (usually the Subject Alternative Name) that uniquely identifies an object in the directory. An IAS server uses this to determine the identity of a user or computer from a presented certificate. IAS uses this implicit mapping rather than looking up certificates in the directory. Certificate publishing and its corollary, certificate lookup, describe where certificates are actually stored in the directory as attributes of user or computer objects. A user can then search for someone in the directory and retrieve certificates belonging to that person.

It is also possible to manually import or map other certificates to user or computer objects using the Active Directory Users and Computers management console.

Conversely, there are many examples in which a direct mapping to a directory object is not required, including the following:

- Web servers where the primary identifier is the DNS host name of the Web site.
- Where certificates are issued to entities that have no Active Directory equivalent (for example, a router or a user from another organization).

All of the end-entities in the solution for this guide have a direct (and implicit) mapping to Active Directory users and computers.

The CAs are special cases in that they are not necessarily mapped to computer objects. However, they are almost always mapped to Certification Authority objects in the Active Directory AIA and Trusted Certification Authority containers.

Creating Certificate Policies

At a minimum, you should define the certificate approval methods for each certificate security category (high, medium, and standard) and Microsoft also recommends doing this for each certificate subject category (computer, internal user, external user). It is unlikely that you will need to define policies at a more detailed level than this. Document these policy decisions as part of your certificate policy statement and CPS.

Different certificate policies are used to indicate the type of certificate approval process and the level of private key security used in the enrollment of a certificate. You can (but do not have to) encode this into the certificate using OIDs to represent the different policies. The rigor of the approvals process should reflect the value of the certificate being issued. The aim of the approvals (or registration) process is to provide a suitable level of confidence that the requester of the certificate is the same entity as the subject of the certificate. The key security strength is a measure of how much confidence you can place in the fact that the private key will remain private and in the sole possession of the certificate subject.

Certificate assurance levels used in this solution were defined earlier in the “Defining Certificate Security Requirements” section of this chapter. You can indicate the assurance level of the certificates that you issue by including a certificate policy OID corresponding to the assurance level in those certificates. Applications (and users) can use the policy to determine how trustworthy a given certificate is.

The three assurance levels defined earlier are mapped onto the three predefined certificate policies in Windows Server 2003 (referred to as Issuance Policies in the Certificate Templates MMC). The following table includes details on how the different certificate policies are used in this solution. You should include the policy OIDs in your certificate templates (at least for medium and high assurance certificates) so that it is easy to identify higher assurance certificates. Certificates with no certificate policy are assumed to be standard assurance.

Table 4.15: Certificate Issuance Policy Levels

Certificate (issuance) policy	Registration requirements	Minimum key storage requirements
Low (Standard)	Automatic approval dependent on successful domain authentication. For stand-alone CAs, this is a level of unauthenticated approval — that is, the CA issues the certificate without performing any (or only a minimal) check on the requester.	Software keys
Medium	Certificate Manager approval. You must define in your policy what types of checks that the certificate manager must perform before he or she approves the certificate request.	Software keys or hardware keys. If using software keys, consider using strong key protection unless the application cannot use this. (For example, computer certificates cannot use strong key protection.)
High	Nominated enrollment officer signature(s) and Certificate Manager approval. You must define the types of checks that the enrollment officer(s) and certificate manager must perform before the request is approved.	Hardware tamper-proof token. For example, smart card or universal serial bus (USB) cryptographic token for a user, or HSM for a computer.

Note: There is no reason why you cannot define more or fewer certificate policies and assurance levels if that makes sense for the requirements of your organization.

Defining Certificate Revocation Policy

In some situations, you might need to invalidate a certificate before it has reached the end of its lifetime. Creating policies for certificate revocation involves the following tasks:

- Defining the conditions that warrant the revocation of a certificate.
- Selecting a CRL publication location.
- Selecting the type or types of CRLs that you intend to use.
- Establishing schedules for the publication of CRLs.

Part of your certificate policy will include defining the conditions that warrant the revocation of the certificate. This may vary for different types of certificates and will very likely vary for certificates of different assurance levels and subject types and for different CAs.

You should also document how revocation reason codes will be used when a certificate is revoked. The following table describes the different reasons codes.

Table 4.16: Certificate Revocation Codes

Reason code	Description
Key compromise	The private key of the certificate has been compromised (or compromise is suspected).
CA compromise	The private key of the CA has been compromised (or compromise is suspected).
Change of affiliation	The subject has moved to a different organization.
Superseded	A new certificate has been issued that takes the place of this certificate.
Cease of operation	The CA is no longer operating.
Certificate hold	Certificate use needs to be temporarily suspended (for example, if a user has misplaced his or her smart card but is unsure whether it is lost).
Unspecified	Any reason not covered by the other codes.

Important: You should avoid using the Certificate hold reason code unless the circumstance is strictly justified. Holding and then releasing a certificate makes it all but impossible to determine the revocation status of a certificate (and therefore the validity of any signature that it makes) at a given time.

As part of the revocation procedure, you should ensure that you document answers to the following questions. These will typically be stored in a change management or incident management log:

- What is the reason for revoking this certificate?
- Who requested the revocation of this certificate?
- Will you ever need this certificate again (such as verification of signatures or decryption of messages)? If so, what is the need (such as verification of signatures, decryption of messages, normal usage)?
- Are there special requirements for the person revoking the certificate (must he or she be a certificate manager for example) that you must fulfill as an administrator?
- Are there any documented operational procedures for my organization that you must follow when revoking certificates (such as backup)?

There are also a number of technical parameters that govern certificate revocation. Microsoft also recommends documenting these as part of your CA policy. The parameters in the following tables are for the root and issuing CAs for this solution.

Table 4.17: Root CA Certificate Revocation Parameters

Parameter	Chosen value	Reason
CRL publication locations (CDPs)	HTTP path to internal Web server	Publishing to the Web server allows backup to LDAP location and allows non-LDAP clients access to the CRL.
	LDAP path to Active Directory CDP container	Publishing to all domain controllers allows easy local access for all domain users.
CRL type	Base CRL only	Due to the small number of certificates ever issued, there is no benefit in using delta CRLs.
Publishing schedule	Six months	This makes revoking issuing CA certificates difficult, so it requires confidence in issuing CA security. However, this long period keeps management of the root CA to a minimum.
CRL overlap period (interval between the new published CRL and the expired CRL)	10 days	This allows some margin of error for retrieving the new CRL from the root CA.

Table 4.18: Issuing CA Certificate Revocation Parameters

Parameter	Chosen value	Reason
CRL publication locations (CDPs)	LDAP path to Active Directory CDP container (for both base and delta CRLs)	Publishing to all domain controllers allows easy local access for all domain users. (See the subsequent note on publishing delta CRLs to Active Directory.)
	HTTP path to internal Web server	Publishing to the Web server allows backup to LDAP location and allows non-LDAP clients access to the CRL.
CRL type	Base CRL Delta CRL	Delta CRLs are useful in optimizing CRL retrieval traffic while giving a relatively short lag for revocation information to be published.
Publishing schedule	Base CRL — seven days	This interval needs to be frequent enough that systems that do not understand delta CRLs still receive relatively fresh revocation information.
	Delta CRL — one day	For modern clients that can use delta CRLs, this provides a relatively small revocation lag.

(continued)

Base CRL overlap period (interval between new CRL being published and old CRL expiring)	Four days	This allows for some margin of error for CA recovery in case it is unable to publish a base CRL on time. Four days was chosen in anticipation of the worst case of a CA failing on a Friday night of a holiday weekend, and no one noticing until the following Tuesday.
Delta CRL overlap period	1 day	Delta CRLs are not service-critical, so it is not catastrophic if a delta CRL publication fails. This is set so that the overlap is greater than the Active Directory replication latency.

Note: Because relatively short-life (one day) delta CRLs are being used, you must ensure that the maximum Active Directory replication latency is less than 50 percent of the delta CRL publication period. Otherwise, this could lead to certificate clients using stale revocation information, as well as possibly having a negative effect on directory replication traffic to sites with constrained network bandwidth. Set the delta CRL overlap time to a value larger than the longest time it takes directory information to replicate through your forest.

If the Active Directory latency is longer than this, or you do not want the extra directory replication traffic, either lengthen the delta CRL publication period or avoid publishing delta CRLs to the directory. If you change the delta CRL locations, you must issue a new base CRL.

Replication latency is typically much less of a concern when using HTTP locations instead of LDAP URLs to publish delta CRLs.

Planning for Key and Data Recovery

Key recovery and data recovery are outside of the scope of this solution. If you require either or both of these, you must plan and manage them carefully to avoid data loss and to prevent inadvertent disclosure of encrypted data.

You should read the sections on "Planning for Data Recovery and Key Recovery" and "Designing a Public Key Infrastructure," in the *Windows Server 2003 Deployment Kit* and the Microsoft technical paper, "Key Archival and Management in Windows Server 2003."

Summary

This chapter covered the process of designing a Public Key Infrastructure (PKI) for a secure Wireless LAN (WLAN). Given the likelihood in many organizations of the PKI being used in the future for other applications, the PKI detailed in this guide is designed with this in mind. The design presented in this chapter is flexible enough to allow you to extend it to cover a wide variety of future requirements. You can use the information here to help you design a PKI that is secure enough to serve much stricter security criteria than those required by the WLAN application.

The design decisions described in this chapter are used in the Build Guide and Operations Guide to implement the PKI. This information is detailed in chapters 7 and 11 of the solution guidance.

In the remaining chapters of the Planning Guide, you will learn about the design of the other core components of this solution — the RADIUS infrastructure (implemented with IAS) and the WLAN security infrastructure.

More Information

The following references provide more detailed background and other related information on many of the subjects covered in this chapter:

- The *Windows Server 2003 Deployment Kit* chapter, “[Designing a Public Key Infrastructure](http://go.microsoft.com/fwlink/?LinkId=4735),” at <http://go.microsoft.com/fwlink/?LinkId=4735>.
- For a general introduction to PKI concepts and the use of Certificate Services in Windows 2000, see [An Introduction to the Windows 2000 Public-Key Infrastructure](http://www.microsoft.com/technet/archive/windows2000serv/evaluate/featfunc/pkiintro.aspx), at www.microsoft.com/technet/archive/windows2000serv/evaluate/featfunc/pkiintro.aspx.
- For a detailed description of the enhanced PKI functionality in Windows Server 2003 and Windows XP PKI, see [PKI Enhancements in Windows XP Professional and Windows Server 2003](http://www.microsoft.com/technet/prodtechnol/winxppro/plan/pkienh.aspx), at www.microsoft.com/technet/prodtechnol/winxppro/plan/pkienh.aspx.
- The Windows 2003 Server Enterprise Edition [Public Key Infrastructure](http://www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/entserver/SE_PKI.asp) product documentation discusses key concepts and administration tasks for Certificate Services at www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/entserver/SE_PKI.asp.
- For more information about writing a certificate practice statement, see RFC2527, [Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework](http://www.ietf.org/rfc/rfc2527.txt) at www.ietf.org/rfc/rfc2527.txt.
- For an example of a CPS, see the [VeriSign Certification Practice Statement \(CPS\)](http://www.verisign.com/repository/CPS/) page at www.verisign.com/repository/CPS/.
- For details on how certificate policy OIDs and CPS URLs are encoded into a certificate, see RFC 3280, [Internet X.509 Public Key Infrastructure Certificate and CRL Profile](http://www.ietf.org/rfc/rfc3280.txt), at www.ietf.org/rfc/rfc3280.txt.
- For a detailed look at qualified subordination, see the technical paper, [Planning and Implementing Cross-Certification and Qualified Subordination Using Windows Server 2003](http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws03qswp.aspx), at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws03qswp.aspx.

- For a detailed explanation of the differences between Enterprise and Stand-alone CAs, see the [Certification Authorities](#) section in the Certificate Services product documentation at www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/entserver/sag_CSCAs.asp.
- For more information about enabling anonymous LDAP access in Windows Server 2003, see the Microsoft Knowledge Base article Q326690, "[Anonymous LDAP operations to Active Directory are disabled on Windows Server 2003 domain controllers](#)" at <http://support.microsoft.com/default.aspx?scid=326690>.
- For a detailed discussion about certificate revocation, see [Troubleshooting Certificate Status and Revocation](#) at www.microsoft.com/technet/prodtechnol/winxpro/support/tshtctrl.msp. The section on CDP Extensions is particularly relevant to some of the discussions in this chapter.
- For more information about installing and using the Certificate Services Web enrollment pages, see the [Product Documentation for Windows Server 2003](#) Web page at www.microsoft.com/windowsserver2003/proddoc/default.msp, and then search for Security, Public Key Infrastructure, Certificate Services, How to..., and Set up a Certification Authority.
- For a detailed guide to Certificate Templates, see [Implementing and Administering Certificate Templates in Windows Server 2003](#) at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws03crtm.msp.
- Descriptions of all of the built-in certificate templates can be found in the Windows Server 2003 product documentation in the [Troubleshooting](#) section at www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/entserver/ctcon_tshoot.asp.
- For more information about certificate auto-enrollment, see [Certificate Autoenrollment in Windows Server 2003](#) at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/autoenro.msp.
- For information about key archival and management in Windows Server 2003, see the technical paper, [Key Archival and Management in Windows Server 2003](#) at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/kyacws03.msp.
- For more information on advanced certificate enrollment scenarios see the paper [Advanced Certificate Enrollment and Management](#) at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/advcert.msp.
- For information on certificate enrollment using the Web enrollment pages see the paper [Configuring and Troubleshooting Windows 2000 and Windows Server 2003 Certificate Services Web Enrollment](#) at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/webenroll.msp.
- For detailed information about implementing a Windows Server 2003 PKI, see the paper "[Best Practices for Implementing a Microsoft Windows Server 2003 Public Key Infrastructure](#)" at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/ws3pkibp.msp.
- Additional implementation information is available in the Microsoft Systems Architecture version 2.0 Implementation Guide, which can be downloaded from the [Windows Server System Reference Architecture](#) page at www.microsoft.com/resources/documentation/msa/2/all/solution/en-us/msa20ik/vmhtm1.msp.

5

Designing a RADIUS Infrastructure for Wireless LAN Security

Introduction

This chapter describes the architecture and design of the Remote Authentication Dial-In User Service (RADIUS) infrastructure used in this wireless local area network (WLAN) solution. The RADIUS infrastructure uses the Microsoft RADIUS implementation—Microsoft® Internet Authentication Service (IAS).

The first goal of the chapter is to describe the design decisions involved in the IAS infrastructure for the solution and to discuss the reasoning behind them.

When phrases like “This solution uses option...” or “This design uses...” appear in this chapter, these refer to decisions made for the solution design that are implemented in the later build and operations chapters of the of the solution guide.

A second goal of this chapter is to help you determine the suitability of the design for your own organization. Therefore, when phrases such as “You should decide this...” appear, these indicate decision points where you need to make a choice based on your own requirements. These decision points generally occur while discussing how you can extend the solution to meet your organization's broader security needs. For this reason, some topics feature a more detailed discussion in order to help you understand the implications of the steps involved, and to prevent you from having to refer to other resources.

Chapter Prerequisites

Before reading this chapter, you should familiarize yourself with RADIUS concepts and IAS deployment options. You can find references to useful resources on these topics in the “More Information” section at the end of this chapter. You will also find useful information in the IAS chapters of the *Microsoft Windows Server™ 2003 Resource Kit*, and the *Microsoft Windows Server™ 2003 Deployment Kit*.

Chapter Overview

The chapter is divided into topic areas that cover the design of a RADIUS infrastructure. The aims of the chapter are to:

- Provide an overview of how you can use IAS to provide a broad network access management solution, and explain how it applies to WLAN in particular.
- Identify the IT environment prerequisites for the solution, and discuss pre-existing infrastructure.
- Detail the design decisions that you face when creating the architecture for an IAS-based RADIUS infrastructure, specifically those decisions that relate to 802.1X-based wireless networking.
- Explore management strategies for maintaining the IAS server infrastructure.
- Provide references to additional information on concepts, product details, and deployment planning.

The following flowchart illustrates the chapter structure.

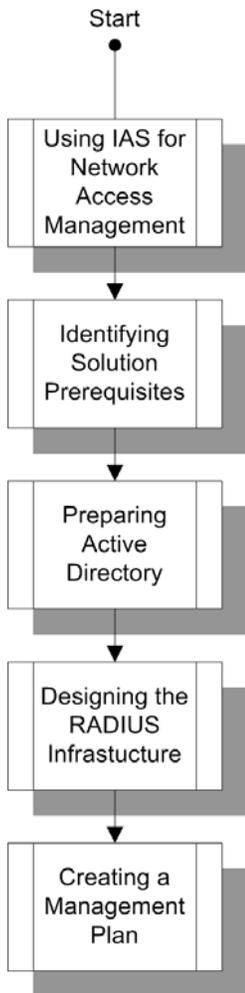


Figure 5.1
Planning an IAS infrastructure

Using IAS for Network Access Management

The Internet Authentication Service in Microsoft Windows Server™ 2003 is the Microsoft implementation of a RADIUS server and proxy server. As a RADIUS server, IAS performs centralized authentication, authorization, and accounting (AAA) of various types of network connections. As a RADIUS proxy server, IAS can forward RADIUS requests to another RADIUS server for AAA. You can use IAS with Virtual Private Networking (VPN) servers, such as the Microsoft Windows®-based Routing and Remote Access Service (RRAS), or with other network access infrastructure, such as wireless access points (AP) and authenticating Ethernet switches.

To maximize the value of an IAS-based RADIUS infrastructure, your organization should make a company-wide decision to use centralized services for network access management. This includes using a centralized accounts database, such as the Active Directory® directory service, and centralizing network access policy management on IAS servers. By centralizing management, the cost of maintaining network access control information on distributed network access equipment can drop dramatically. In addition, centralizing accounts and network access policy helps to reduce the security risks associated with configuring and managing distributed equipment.

Planning and deploying an IAS infrastructure that meets your organization's current and future needs requires careful consideration. IAS is not intended to provide access to a single, isolated network. Instead, you should deploy IAS to provide strategic network access management for a variety of network access scenarios.

Identifying Your Organizational Network Access Management Requirements

IAS in Windows Server 2003 supports a number of network access scenarios, including:

- **Wireless access.** You can configure wireless access points capable of 802.1X to use IAS AAA services for 802.11-based WLAN access control, and to provide key management.
- **Wired access.** Ethernet switches capable of 802.1X can use IAS AAA services for per-port access control to wired LANs.
- **VPN access.** VPN servers such as Windows-based RRAS can use IAS AAA services for corporate network access control, as well as to provide key management.
- **Dial-up access.** Dial-up servers such as Windows-based Routing and Remote Access can use IAS AAA for network access control for the corporate network.
- **Extranet access.** Extranet access servers can use IAS AAA services when providing restricted access for business partners to shared resources.
- **Outsourced corporate network access.** Network solution providers can exploit IAS AAA services to integrate outsourced network infrastructure with customers' account databases and access control policy. IAS also can provide accounting information required to bill customers for this service.
- **Internet access.** Internet Service Providers (ISP) can exploit IAS AAA services to provide dial-up and high speed Internet access while using individual organizational account databases and access control policy. IAS can provide accounting information required to bill customers for this service.

To maximize your investment in IAS and minimize future change to your IAS infrastructure, you should evaluate each of these scenarios for use in your organization. Although IAS in this solution is only used for wireless network access, you can extend the solution to support these other scenarios. Information on extending RADIUS infrastructure to support additional scenarios is available in Chapter 3, “Secure Wireless LAN Solution Architecture.”

Using IAS for Wireless Network Access Management

WLANs are becoming more widespread with the advent of industry standards such as Institute of Electrical and Electronics Engineers (IEEE) 802.11 standards. WLANs allow a user to roam around a building or campus and automatically connect to the network as long as the user is in the proximity of a wireless AP.

While providing convenience, WLANs present the following security risks:

- Anyone who has a compatible WLAN adapter can potentially gain access to the network.
- Wireless networking signals use radio waves to send and receive information. Anyone within an appropriate distance to a wireless AP can detect and receive all data sent to and from the wireless AP.

To counter the first security risk, you can set up wireless APs as RADIUS clients, and then configure them to send access requests and accounting messages to a central RADIUS server running IAS. To counter the second security risk, you can encrypt the data sent between the wireless devices and the wireless APs.

IAS provides enhanced security for WLANs in two ways: it performs as a RADIUS server for IEEE 802.1X wireless APs and client devices; and it provides dynamic encryption keys through certificate based authentication protocols, such as the Extensible Authentication Protocol–Transport Layer Security (EAP–TLS) protocol.

Note: Throughout this guidance, wireless APs may be referred to as RADIUS clients. Although wireless APs are not the only type of RADIUS client possible, they are the only RADIUS client type addressed in this guidance. For this reason, the two terms are used interchangeably.

Identifying Prerequisites for the Solution

You should be sure to understand what pre-existing conditions are required in your environment before starting to design a wireless access management solution using IAS.

Active Directory Considerations

This solution is designed for organizations that have deployed Active Directory and run Windows 2000 Server or later on their domain controllers. These are prerequisites in this solution because several RADIUS design decisions have been made that require features that are only available in domains using Windows 2000 native mode or higher. The following table illustrates some features used in this solution and their level of support with various domain functionality levels.

Table 5.1: Windows Domain Features Leveraged in the Solution

Feature	Windows Server 2003 native mode	Windows 2000 native mode	Mixed mode -or- Microsoft Windows NT® 4.0
Universal and nested groups	Yes	Yes	No
User principal names (UPN)	Yes	Yes	No
Control access through Remote Access Policy (RAP) permission available in user account	Yes	Yes	No
Support for EAP-TLS	Yes	Yes	No

Note: The Certificate Services implementation for this solution also has requirements specific to Active Directory. For more information, see Chapter 4, “Designing the Public Key Infrastructure.”

Although it is not required, after reading this chapter you may decide to deploy IAS on domain controllers. This solution is based on IAS in Window Server 2003 and for this reason would require you to upgrade the target domain controllers to this operating system version. More information about co-locating IAS with domain controllers is provided later in this chapter.

Pre-existing RADIUS Infrastructure

This solution makes no provision for integrating with existing RADIUS servers in your environment. However, existing IAS-based and third-party RADIUS servers can integrate with this solution. In most cases you will want to use Windows Server 2003 IAS for features related to WLAN access.

You can upgrade older versions of Windows RADIUS servers to Windows Server 2003 to serve as the core RADIUS servers in this solution. Alternatively, you can modify the existing RADIUS servers to proxy RADIUS traffic to the new Windows Server 2003–based RADIUS servers.

For detailed planning guidance around migrating existing RADIUS infrastructure to Windows Server 2003 IAS, see your Microsoft partner or contact your Microsoft Account Executive who can connect you with the appropriate partner or Microsoft Consulting Services professionals.

Designing the RADIUS Infrastructure

You must make a number of design decisions when using IAS to support 802.1X-based WLAN access. This section describes several of these decisions and discusses the options selected for this solution. You should evaluate each decision in terms of how it applies to your environment.

Determining the Role of IAS as a RADIUS Server

You can deploy IAS servers to function in one of three conceptual RADIUS roles:

- RADIUS server
- RADIUS proxy server
- RADIUS server and proxy server

Note: Throughout this guidance, the terms RADIUS server and RADIUS proxy server are used to describe an IAS server configured to perform these roles.

The following table details some of the capabilities of servers configured to perform these roles and identifies how they are useful in real-world scenarios.

Table 5.2: IAS RADIUS Roles

IAS RADIUS role	Capabilities	Scenario
RADIUS server	<ul style="list-style-type: none"> – Checks credentials directly against Active Directory or other authoritative data sources. – Leverages RAP to determine network access. 	Required for all network access management scenarios.
RADIUS proxy server	<ul style="list-style-type: none"> – Routes request based on request properties. – Can modify the RADIUS properties of requests in transit. – Provides load balancing of RADIUS requests to RADIUS Server groups. 	<ul style="list-style-type: none"> – Useful in multiforest scenarios where network access equipment is shared. – Useful for deploying high-scale front-end and back-end network AAA architectures. – Useful for federating authentication to external organizations.
RADIUS server and proxy server	Combination of the previous two capabilities.	Combination of the two scenarios.

Not all RADIUS roles are required for all network access management scenarios. For example, WLAN access management will only require the RADIUS server role for many organizations. However, if your organization plans to use wireless network infrastructure to service users and devices from multiple Active Directory forests, you will also need a RADIUS proxy server role to route requests to separate RADIUS servers in each forest.

For simplicity and cost reasons, this solution only includes IAS servers configured as RADIUS servers. It does not implement IAS as a RADIUS proxy server.

Understanding Server Failover and Load Balancing

RADIUS is a critical component of any 802.1X-based WLAN access management solution. The availability of IAS servers to wireless APs determines the availability of the WLAN to end users. Therefore, you should take care to ensure that two or more IAS servers are available to wireless APs at all times. Most modern wireless APs include the ability to configure two RADIUS servers for authentication, and two RADIUS servers for accounting. This ensures that the loss of contact to a single RADIUS server does not affect service to the WLAN clients.

To implement multiple servers for resiliency, many organizations will want to select a scheme to load balance the requests from the wireless APs that are configured as RADIUS clients across the RADIUS servers to ensure that none of the servers is constrained.

Before choosing a load balancing strategy, it is important to understand that 802.1X implements EAP within RADIUS (EAP–RADIUS) between the wireless APs and the RADIUS servers. Although RADIUS uses the connectionless User Datagram Protocol (UDP), EAP is a session-oriented protocol that is tunneled within RADIUS. This effectively means that multiple EAP–RADIUS packets that comprise a single authentication operation must return to the same RADIUS server or the authentication attempts will fail.

The following table illustrates several options to ensure that the RADIUS clients can use both multiple RADIUS servers for resiliency and for load balancing RADIUS requests.

Table 5.3: EAP–RADIUS Failover and Load Balancing Options

Failover and load balancing method	Advantages	Disadvantages
IAS proxies with RADIUS Server Groups	<ul style="list-style-type: none"> – RADIUS service failure detection with failover and failback. – Distribution of traffic load based on traffic properties. – Maintains EAP session state when load balancing. – Configurable request distribution to servers based upon priority and weight settings. 	<ul style="list-style-type: none"> – Additional IAS servers required. – Still requires APs to be configured with primary and secondary proxy RADIUS IPs.
Primary and Secondary RADIUS server settings on wireless APs	<ul style="list-style-type: none"> – Simpler configuration for small environments. – Wireless AP detects traffic failure and performs failover. – Uses native wireless AP functionality. 	<ul style="list-style-type: none"> – Requires careful planning and monitoring of primary and secondary RADIUS server selection. – Many wireless APs do not support failback functionality leading to unbalanced load servers.

Enterprise organizations and large network service providers should consider using RADIUS proxies to accept requests from RADIUS clients and distribute the load to RADIUS servers, which you can configure into RADIUS Server Groups. You can base the distribution of network traffic to the RADIUS servers in the RADIUS Server Groups on a number of configurable items. These items include RADIUS traffic type and RADIUS attributes, in addition to priority and weighting values. The RADIUS servers in each RADIUS Server Group can then perform core authentication and authorization for users and devices within a domain or an entire forest. This creates a front-end, back-end architecture for servicing RADIUS requests, and provides the most flexibility for load balancing and scaling options.

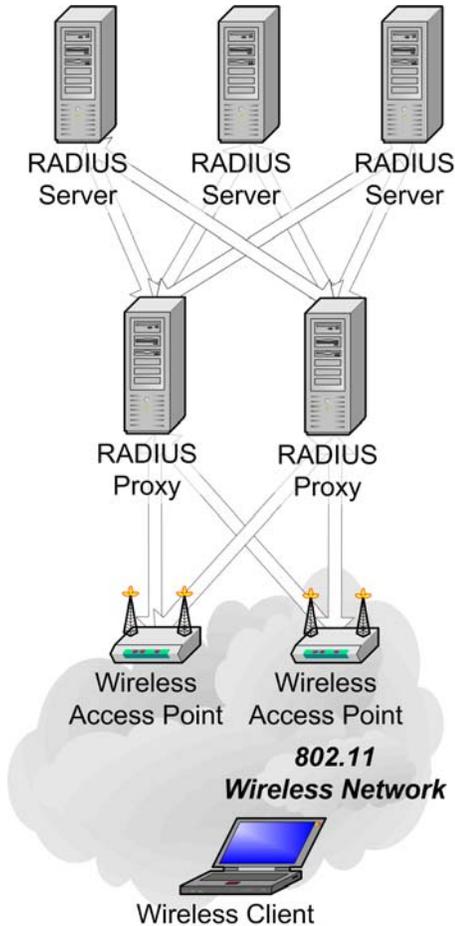


Figure 5.2

Failover and load-balancing using RADIUS proxies

However, the simple RADIUS server failover capabilities in modern wireless APs provide a level of resilience that is sufficient for most organizations. If this is not sophisticated enough, the migration path from this to a RADIUS proxy server-based failover and load-balancing strategy is relatively straightforward. A disadvantage of using a wireless AP-based failover and load-balancing strategy is the management overhead required to pair wireless APs with RADIUS servers, monitor the RADIUS servers for uneven service loads, and make modifications as needed. Another disadvantage is that some models of wireless AP do not support failback. Failback occurs when an AP that has failed over to use a secondary RADIUS server automatically switches back to its designated primary RADIUS server as soon as the primary server has recovered. Without failback, all of the

wireless APs might fail over to their secondary RADIUS server and then require administrator intervention to repoint them to the correct primary server.

To achieve load balancing using a wireless AP-based failover strategy when multiple RADIUS servers are available locally:

- Configure half of the wireless APs in each location to use the primary RADIUS server first, and the secondary RADIUS server in the event that the primary server fails.
- Configure the other half of the wireless APs in each location to use the secondary RADIUS server first, and the primary RADIUS server in the event that the secondary server fails.

Note: The terms “primary” and “secondary” do not denote any difference in functionality between the servers—they are peers. These terms are used here to distinguish between the servers for the purposes of the failover discussion.

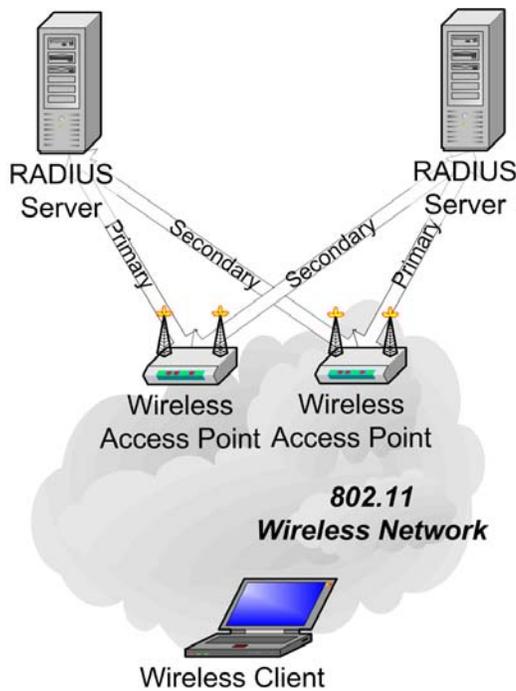


Figure 5.3
Wireless AP-based failover and load balancing

To achieve load balancing using a wireless AP-based failover strategy in branch office situations when one RADIUS server is available locally but a remote RADIUS server is also available, you should configure all of the wireless APs in the branch office to use the local RADIUS server as the primary server. Then configure the remote RADIUS server as the secondary server for use upon primary server failure.

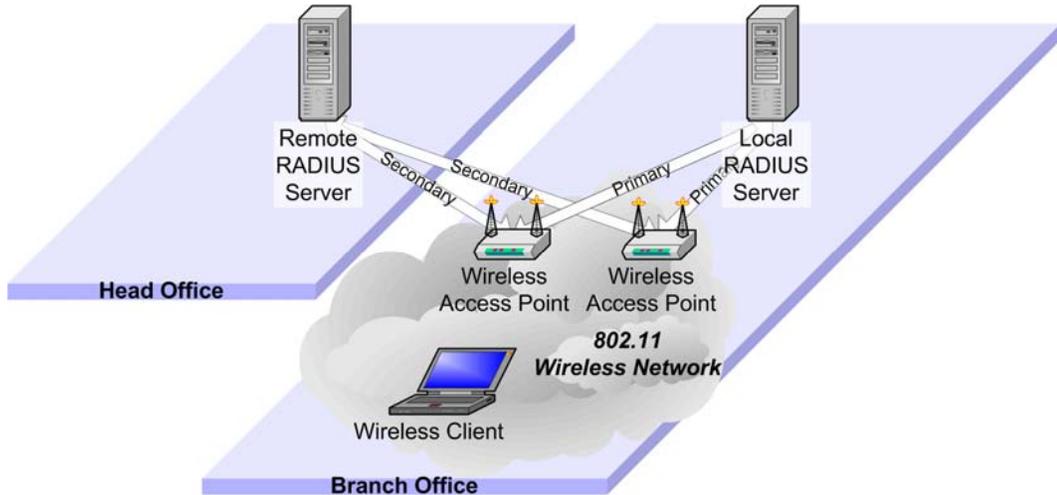


Figure 5.4

Wireless AP-based failover and load balancing with local and remote RADIUS servers

For branch office scenarios, ensure that wireless APs have the capability to fail back to the primary RADIUS server when it returns to service. Otherwise, you must manually reconfigure the wireless APs to avoid unnecessary wide area network (WAN) traversal for the RADIUS service.

Note: Ask your hardware vendor about the support for fail-back in its products.

Branch office RADIUS servers are optional; you can use centralized RADIUS servers across a WAN instead. However, without a local RADIUS server and domain controller, users in the remote office cannot access the local WLAN if the WAN fails.

This solution is designed to use wireless AP-based server failover and manual configuration for load balancing. For more information about planning a RADIUS infrastructure that uses RADIUS proxies for server failover and load balancing, see the “Deploying IAS” chapter of the *Microsoft Windows Server 2003 Deployment Kit*. You can find a reference to this resource at the end of this chapter.

Establishing Logging Requirements

You can configure IAS servers to log two types of optional information:

- Successful and rejected authentication events.
- RADIUS authentication and account information.

Successful and rejected authentication events generated from devices and users attempting to access the WLAN are recorded by IAS in the Windows Server 2003 System Event Log by default. The authentication Event Log information is most useful for troubleshooting authentication issues, although you also can use this information for security auditing and alerting purposes.

You should leave the options for **success** and **reject** event logging enabled initially, but you may want to disable the **success** option once the system has stabilized. This is because successful WLAN access events will rapidly fill the System Event Log, and using this option also may be unnecessary for security purposes if the option for RADIUS authentication request logging is enabled.

Enterprise-level organizations should consider using enterprise monitoring tools such as Microsoft Operations Manager (MOM) to act on IAS events in the System Event Log using custom scripts. For example, a custom MOM script could detect an increase in IAS events related to rejected authentication attempts, and then notify an administrator to take action.

IAS also offers the ability to save authentication and network access session information in the form of RADIUS request logs. You can selectively enable and disable options to provide the following information from your RADIUS request logs:

- Accounting requests—For example, accounting start and stop messages that indicate the start and end of a network access session.
- Authentication requests—For example, access-accept or access-reject messages that indicate success or failure of authentication attempts.
- Periodic status—For example, interim accounting requests that some network access devices send.

RADIUS request logs are most useful for organizations such as network service providers that charge customers a fee based on usage of the network. However, you also can use RADIUS request logs for security monitoring and auditing. In particular, RADIUS authentication and accounting logs allow security auditors to determine such things as:

- The details of unauthorized authentication attempts to the WLAN.
- The durations of accepted connections to the WLAN.

IAS can log to text logs or Microsoft SQL Server™ 2000 databases. Text based logging of RADIUS authentication and accounting information is disabled by default in IAS. Before enabling RADIUS text-based logging you should:

- Speak to your security staff to understand the requirements to track WLAN access information, and which details are required.
- Perform lab testing of RADIUS text logging to understand the server hardware requirements (disk and CPU) during load balancing from your WLAN users. WLAN access can generate significantly more information than other types of network access types.
- Evaluate which RADIUS request information (authentication, accounting, and periodic status) is required and which is optional. WLAN access can generate significant amounts of information that can rapidly consume disk space.
- Determine a strategy for accessing, storing, and archiving RADIUS request log information. You can save RADIUS request log information as text files on the hard disk of each IAS server or in a SQL Server database.

Enterprise-level organizations that need to use RADIUS accounting logs will want to consider using the SQL Server logging features of IAS. You can log the RADIUS accounting information to the SQL Server Desktop Engine (MSDE) on each IAS server and then replicate the information to a central SQL Server cluster. This strategy provides centralized and structured storage of RADIUS accounting data to facilitate queries, reporting, and archiving. Performing SQL Server logging to local MSDE databases also eliminates the possibility of network issues preventing IAS from logging this information, and network access requests that rely on the account information from being rejected.

Organizations without SQL Server 2000 or staff to perform regular queries, reporting, and archiving of RADIUS request logs will still want to consider recording this information for security incident investigations. Several design decisions related to RADIUS logging have been made in this solution as indicated in the following table. Review this information to determine which ones comply with the requirements of your environment.

Table 5.4: IAS Logging Design Decisions

IAS logging design decision	Comments
The System Event Log size in the IAS group policy template in the <i>Windows Server 2003 Security Guide</i> was increased from the defaults to accommodate IAS events.	If you choose not to enable the option for RADIUS authentication request logging, then the System Event Log is the primary record of WLAN access security events by default. Carefully consider settings such as the default Overwrite events as needed setting because this may allow audit data to be overwritten once the log is full.
RADIUS authentication and accounting request logging to text files was enabled.	This introduces CPU load and disk space requirements on the IAS servers. IAS will stop accepting authentication and accounting requests if logging cannot be performed. For this reason, consider the potential of denial of service (DoS) attacks to fill up the log file disk.
The IAS server hardware specifications in this guidance include a separate log file disk volume on separate physical disks.	This ensures that write performance of RADIUS request log files has minimal performance impact on RADIUS network access management. This decision also ensures that events that cause logging to fill a disk volume do not affect the ability of the server to recover.
The options for RADIUS authentication and accounting items were enabled, but the option for periodic status was not.	This ensures that only essential information required to determine authentication status and session duration are logged. The option for periodic status has been omitted to reduce log file requirements. You should evaluate enabling periodic status logging if recording user session durations is important in your environment.
Open Database Connectivity (ODBC)-compatible database format was chosen for the RADIUS authentication and accounting log files.	This allows administrators to easily import log files into ODBC-compliant databases for analysis and may generally be considered a best practice. In addition, you can use IASPARSE.exe from the Windows Server 2003 Support Tools to browse the files.
The interval for creating new log files is set to Monthly .	Selecting an interval that produces fewer log files facilitates importing the log files into databases or using IASPARSE.exe to browse the log files if you do not use SQL Server logging. Weigh this option against the risk of filling a hard disk with a single log file.
An option in RADIUS request logging was set to delete the oldest log file first when the disk is full.	The risk of this setting (the default) is that security information may be lost in the event that a log file disk becomes full. This setting has been chosen to avoid stopping the IAS servers if the log file become full. If preserving the security logs is more important than service availability, you should disable this setting.

Choosing to Centralize or Distribute Servers

The decision to use centralized or distributed IAS servers is based in part upon the geographic distribution of your organization and your organization's IT infrastructure deployment strategy. Consider which of the following three IT infrastructure strategy types your organization most closely follows:

- A centralized IT infrastructure
- A distributed IT infrastructure
- A mixed IT infrastructure

Many modern IT organizations strive to provide fewer, more fault resilient, and more centralized IT infrastructure components. To achieve this goal requires significant investment in high-speed and fault-resilient WAN infrastructure to ensure that branch office users receive the same level of IT service as centrally located users. One advantage of this strategy is that you can redirect the cost of distributed server infrastructure to network infrastructure and bandwidth. In addition, the server infrastructure is located closest to trained data center operations and engineering staff, thus achieving a higher level of availability.

Centralizing IAS servers in organizations that have high-speed, resilient WANs can help reduce the cost of an 802.1X WLAN solution. This type of IT infrastructure strategy and should be considered the starting point of a RADIUS server design for enterprise-level organizations. The RADIUS protocol does not consume a lot of network bandwidth and works well over WAN links. You must also consider the possibility of protocols such as Dynamic Host Configuration Protocol (DHCP) that may time out while waiting for 802.1X authentication to complete. In addition, it is critical to have a high performance connection between the IAS servers and domain controllers that contain the users and groups your environment uses to determine network access. You can avoid many potential issues with 802.1X networking by ensuring that high speed communications between IAS servers and Active Directory are maintained.

However, for some IT organizations the cost of bandwidth, sophisticated network equipment, and redundant WAN connections prohibits them from following centralized IT infrastructure model. These organizations will instead choose to follow a decentralized IT infrastructure model with a server infrastructure distributed to branch offices. This model will ensure continued IT service in the event of a WAN failure.

A third type of IT infrastructure strategy exists that allows organizations to centralize their IT infrastructure when possible, and distribute their IT infrastructure when required. This strategy allows for grouping most IT infrastructure in hub locations to service hub location users, and branch office users connected to the hub. At the same time, this model allows you to distribute server infrastructure to branch offices with many end users. The following diagram illustrates an example of such a mixed IT infrastructure organization.

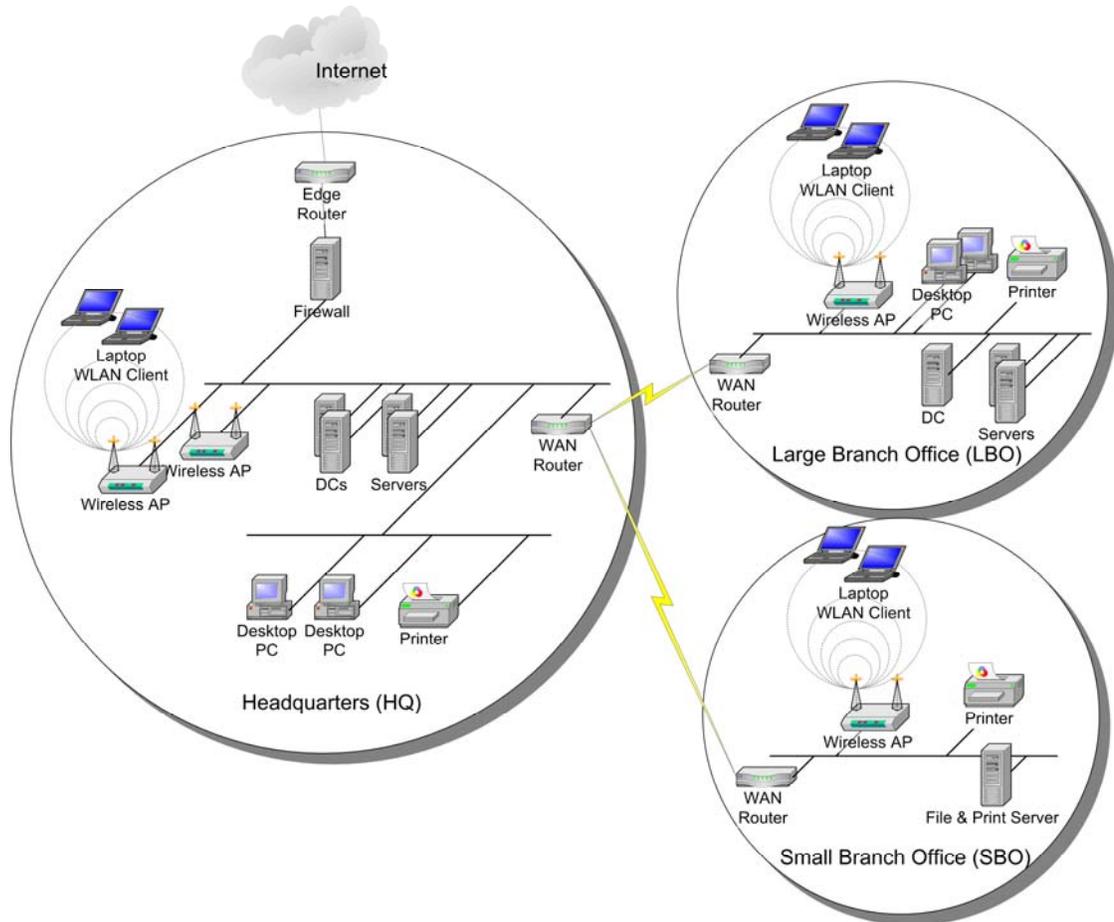


Figure 5.5

An organization with a mixed IT infrastructure that is both centralized and distributed

The solution in this guidance has been designed to accommodate the centralized, decentralized and mixed server infrastructure deployment models by providing the following:

- Guidance for configuring large hub offices with two RADIUS servers that can service local requests and other requests from offices without a server infrastructure.
- Guidance for configuring large branch offices with an optional branch office RADIUS server.

Note: For small branch offices without a server infrastructure, access to WLANs depends on WAN availability.

Determining the Number and Location of Servers

At a minimum, each independent Active Directory forest should have at least two IAS servers acting as RADIUS servers for forest users and devices. This ensures that network access requests will continue to be serviced if one of the RADIUS servers becomes unavailable.

Head office locations with many users are good candidates for two or more RADIUS servers. If high-speed bandwidth is available between multiple hub locations with RADIUS servers, you can configure your wireless APs to fail over to the RADIUS servers located across a WAN. However, when planning to use RADIUS servers across a WAN, assess whether you have an adequate network link between the RADIUS servers and the domain controllers that they depend on in your environment. You should also test timeout values on the wireless APs and client computers, and modify the AP parameters as required. Finally, locate the RADIUS servers in the root domain of your forest to optimize Kerberos operations.

Branch offices that are large enough to warrant domain controllers and do not have resilient WAN connections to hub locations are candidates for a local RADIUS server. If your organization does not have WAN resiliency, weigh the initial and ongoing cost of a branch office IAS server against the cost of wireless users not having access to the wireless network when the WAN is not available.

Determining Co-Location of IAS with Other Services

Due to the intensive communications between IAS and Active Directory domain controllers, you can achieve a performance gain by running IAS on the same server as your domain controllers (this avoids latency issues related to communicating over the network). However, you should carefully consider the implications of co-locating IAS on your domain controllers. The following table details some of these considerations.

Table 5.5: IAS and Domain Controller Co-Location Considerations

IAS location	Advantages	Disadvantages
Co-located on domain controllers	<ul style="list-style-type: none"> – Performance increase for user and computer authentication and authorization. – Requires less server hardware. 	<ul style="list-style-type: none"> – No separation of IAS administrators from domain administrators – No inherent separation of fault or performance issues of co-located services.
Separate from domain controllers	<ul style="list-style-type: none"> – Separation of IAS administrators from domain administrators. – IAS load and behavior does not affect Active Directory service. 	Requires extra server hardware

The Active Directory domain controllers are critical IT infrastructure that you should treat with great care. Many enterprise organizations have a policy of limiting additional software or services on domain controllers to ensure their maximum service reliability.

In many enterprise-level organizations, the RADIUS administrators have separate job roles from Active Directory administrators. IAS is an optional component of Windows operating system, and there is no inherent separation of IAS administration from responsibilities that Windows Local Administrators carry out. For this reason, when IAS is installed on the domain controllers, the IAS administrators are members of the Domain Administrators Security Group.

This solution requires the Windows Server 2003 version of IAS. For this reason, you must upgrade your domain controllers to Windows Server 2003 (if you have not already done this). Consider the following prerequisites prior to upgrading domain controllers in your environment running Windows 2000 Server to Windows Server 2003.

Table 5.6: Prerequisites for Windows Server 2003 Domain Controllers

Issue	Prerequisite	Comment
Windows Server 2003 domain controllers require Server Message Block (SMB) signing and encryption or signing secure channel communications by default. This requirement can cause issues with some earlier versions of Windows-based clients.	Upgrade all of the client computers in your environment to at least Microsoft Windows® 95 with the Active Directory client or Windows NT 4.0 with Service Pack 4 (SP4) or later.	See the <i>Windows Server 2003 Help and Support Center</i> for more details referenced in the More Information section at the end of this chapter.
Windows Server 2003 domain controllers require Secure Channel signing and encrypting by default. This requirement can interfere with domain trusts to servers in domains running Windows NT 4.0 without SP4.	Upgrade all the domain controllers in the legacy domain to Windows NT Server 4.0 with SP4 or later.	See <i>Windows Server 2003 Help and Support Center</i> for more details referenced in the More Information section at the end of this chapter.
Windows Server 2003 domain controllers require Active Directory forest and domain preparation prior to installation.	Prepare the new forest by using the ADPrep utility prior to upgrading the domain controllers in your environment to Windows Server 2003.	This does not affect the Partial Attribute Set (PAS) and thus does not cause a Global Catalog server rebuild.

This solution has been architected to allow for the co-location of IAS and Active Directory on your domain controllers if desired. The solution was tested with IAS separate from Windows Server 2003 domain controllers in hub locations, and co-located with Windows Server 2003 domain controllers in branch offices.

Estimating RADIUS Server Load

IAS performs well on modest server hardware and can be scaled up using additional hardware or scaled out by using RADIUS Server Groups. However, it is best to estimate

the load that WLAN clients will incur on IAS server hardware in advance to avoid server resource constraints that can affect availability of the service.

An optimal design should include the minimum number of servers required for resilience, while leaving room for future growth. Leaving room for growth is particularly important when choosing server hardware for use in a wireless AP-based load balancing model. Moving from wireless AP-based load balancing to RADIUS proxy server-based load balancing, the number of servers needed can jump from two to five (assuming that existing RADIUS servers have reached maximum capacity).

IAS server load considerations include:

- The number of users and devices requiring authentication and accounting.
- The authentication options such as EAP (Extensible Authentication Protocol) type and reauthentication frequency.
- RADIUS options such as logging and IAS software tracing.

Estimating the number of users and devices requiring WLAN access is necessary for estimating IAS server load. Some organizations limit the use of WLANs to a subset of their user population such as executives, whereas other organizations may choose to offer WLAN access to all users. Regardless of which strategy your organization chooses, you must estimate the 'worst case' scenario: where all of your WLAN-enabled users and devices require authentication and authorization within a short period of time. This ensures that IAS servers will be sized to accommodate periods of high stress such as peak office hours and shortly after a major network outage.

The WLAN authentication options chosen have a large effect on IAS server load. Certificate based protocols such as EAP-TLS perform a CPU intensive public key operation upon initial logon but then use a cached credentials strategy called fast reconnect for each subsequent logon until the cache expires (eight hours by default). Full reauthentication will always occur when a wireless client switches from an AP authenticating to one IAS server to an AP authenticating to a different IAS server (for example, when the client moves through a building). This roaming reauthentication only happens once between each client and IAS server and is transparent to the end user when EAP-TLS is used.

You can force wireless clients to reauthenticate to the RADIUS servers as a method to refresh 802.11 WEP session encryption keys. Some wireless AP models include features to perform timed WEP session key refresh without the RADIUS server having to force clients to perform periodic reauthentication. This type of feature is vendor specific. The WiFi Protected Access (WPA) standard includes enhanced encryption and key management features that alleviate the need for forced reauthentication to refresh the session keys.

Therefore, when you are building a model for how many authentications each IAS server will service, consider the different types of authentication behavior in the following table.

Table 5.7: EAP–TLS Authentication Behavior

Authentication type	Comments
Initial computer authentication	Client performs a full authentication with IAS.
Initial user authentication	Client performs a full authentication with IAS.
User reauthentication when roaming between wireless APs	Client performs a full authentication once with each IAS server then uses fast reconnect for additional authentications.
Device reauthentication when roaming between wireless APs	Client performs a full authentication once with each IAS server then uses fast reconnect for additional authentications.
Timed computer reauthentication	Client uses a cached authentication with IAS.
Timed user reauthentication	Client uses a cached authentication with IAS.

Estimates on the number of authentications that IAS can service are best represented as authentications per second. IAS can achieve the following numbers on a computer running Windows Server 2003 with Active Directory that uses an Intel Pentium 4.2–gigahertz (GHz) CPU.

Important: The information in the following table is provided without warranty of any kind and should only be used as a guideline for capacity planning purposes, not for performance comparisons.

Table 5.8: Authentications Per Second

Authentication type	Authentications per second
New EAP–TLS authentications	36
New EAP–TLS authentications with offload card support	50
Authentications with fast reconnect	166

You can configure IAS to generate disk-based text logs containing varying amounts of RADIUS request information. Plan to use a high performance disk for storing the RADIUS logs because of the overhead that RADIUS logging has on the RADIUS servers. Slow disk subsystems can delay IAS RADIUS responses to wireless APs, leading to protocol timeouts and unnecessary failover of wireless APs to secondary RADIUS servers.

In addition, enabling Windows 2003 Server software tracing features will apply an additional load on your IAS servers. However, this may be required occasionally to troubleshoot network access issues. For these reasons, scale your IAS servers with the capacity to run with the tracing features enabled for limited periods while continuing to handle the production load.

Estimating Server Hardware Requirements

Select the server hardware for IAS from the [Windows Server 2003 Hardware Compatibility List \(HCL\)](#). Selecting server hardware from the Windows Server 2003 HCL will help you to avoid reliability and compatibility issues that may occur with untested hardware and device drivers.

Ensure that your IAS servers meet the recommended hardware requirements for Windows Server 2003. And be sure to take into account other services that may also be running on the system, such as Active Directory. Consider using IAS server hardware that can scale to double your expected per-server authentication load. Scaling the capacity of your servers this way will ensure that the appropriate server resources are available to handle server failover scenarios and unusual network conditions.

Using the server layout for Woodgrove Bank, a fictitious company, the following figure displays an example of the IAS-based RADIUS server design. The figure displays the location and number of IAS servers required for the company's expected user distribution and load. However, only a subset of this infrastructure was tested for the guide (using the London hub and the Johannesburg regional offices).

Note: Woodgrove Bank is a fictitious company that represents a medium to large organization. Its network architecture and characteristics were used as the basis for a variety of design decisions for the implementation of this solution.

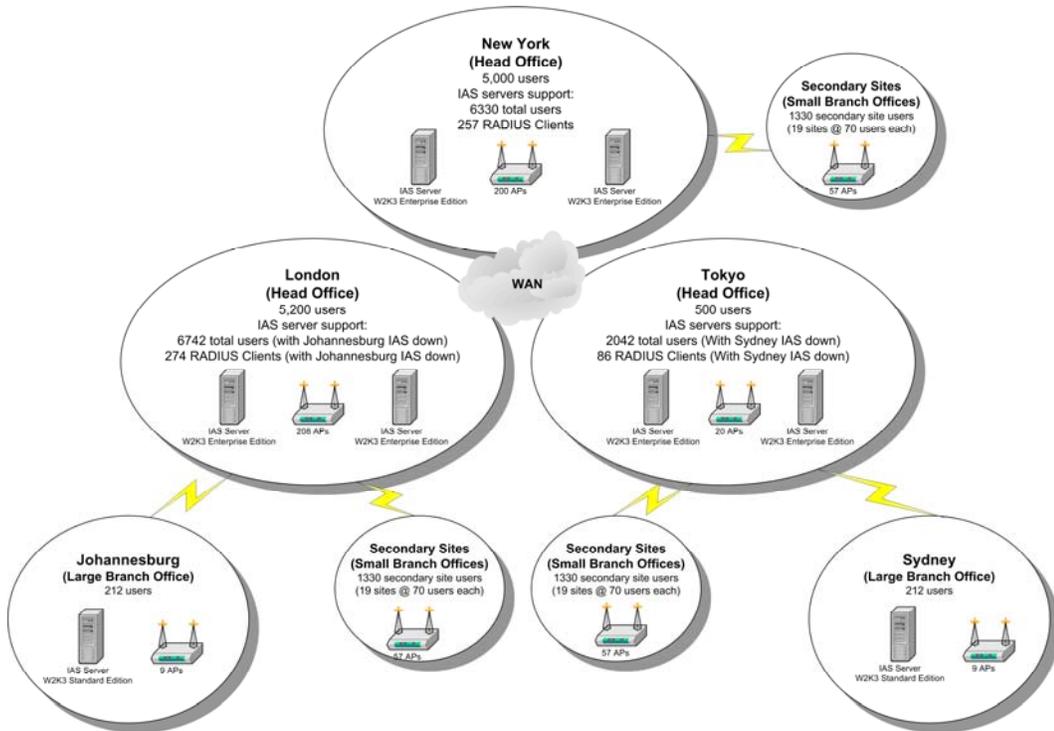


Figure 5.6
The user, wireless AP, and IAS server distribution for Woodgrove Bank

This solution was designed with two RADIUS servers located in a head office housing 6,742 users. Only 50 percent of the users were wireless enabled; thus, 3,371 users and their 3,371 assigned devices can authenticate using EAP–TLS to the two RADIUS servers during peak load periods. Each server can scale to service 3,371 authentications during a 30-minute peak logon period. The peak logon period load equates to roughly two new EAP–TLS authentications per second with the capacity to complete four new authentications per second during times of server failover.

The server was configured to log RADIUS Authentication and accounting requests to text files. This server was a dedicated RADIUS server; domain controller services were located on other servers. Excess capacity was intentionally built into the RADIUS server hardware specification to accommodate possible future network access control requirements for applications such as VPN, wired networks, and dial-up access.

The following table details the IAS server hardware used during the testing of the solution.

Table 5.9: Server Hardware Tested

Resource	Configuration
CPU	Dual CPU Pentium III 850 MHz
RAM	512 MB (megabytes)
Network interface card (NIC)	Two NICs teamed for resilience
Hard disk	– Two, 9 GB hard disks in a RAID–1 configuration (volume C) for the operating system – Two 18 GB hard disks in a RAID–1 configuration (volume D) for log files and configuration data

The hardware requirements for your IAS servers will likely be different. Evaluate your requirements based on the specific variables of your organization.

Determining Server Software Requirements

You must determine whether you need to use Windows Server 2003 Standard Edition or Enterprise Edition for the IAS servers in your environment. Windows Server 2003 Standard Edition is limited to supporting 50 RADIUS clients (Wireless APs, for example) and two Server Routing Groups.

This solution was tested with Windows Server 2003 Enterprise Edition for the hub office RADIUS servers, and Windows Server 2003 Standard Edition for the branch office RADIUS servers. However, the solution will work equally well with the previously mentioned limitations of either software version.

There will be other software components required for your environment depending on the standards in your organization, such as the following:

- Backup agents.
- Management agents such as MOM or Microsoft Systems Management Server (SMS) client components.
- Antivirus software.
- Intrusion detection agents.

Creating a Management Plan

IAS-based RADIUS servers require relatively little ongoing maintenance to ensure continued service availability and network security. However, you should determine your IAS management strategy at the start of your WLAN project so that you can ensure that the appropriate staff are trained and equipped to manage the RADIUS infrastructure.

Change and Configuration Management

Maintaining a known state on your IAS servers is essential to ensuring service availability and network security. IAS natively facilitates transactional changes to various server configuration elements via the **netsh** command, and thus allows for easy rollback in the event that a change causes unplanned behavior.

The **netsh** command allows you to export and import all or portions of IAS configuration to text files. You can use these files to replicate settings between your IAS servers. This feature can speed deployment of configuration changes in large environments.

The tasks required for appropriate change and configuration management are listed in Chapter 12, “Managing Your RADIUS and WLAN Security Infrastructure.”

Planning for Service Recovery

To ensure the rapid recovery of your RADIUS service in the event of a disaster requires careful planning prior to the event. You can streamline the installation and configuration of IAS by using the installation scripts provided with this guidance and you can easily automate the steps required to rapidly restore the IAS configuration state using **netsh** scripts. You will find more information about service recovery tasks in Chapter 12, “Managing Your RADIUS and WLAN Security Infrastructure.”

Planning Administrative Permissions

IAS is an optional component of the Windows Server 2003 operating system and thus does not require a separate administrative security model from the one used on the local server. Complete separation of IAS administration from local server Administrators is not possible. You can achieve some separation without custom development, such as creating a secured Web application that makes IAS configuration changes using a privileged account with local server administration permissions.

However, it is still important to plan the types of administration required and the access requirements to various IAS resources to achieve a least privilege-based model. The following table includes examples of roles and tasks that are related to IAS servers.

Table 5.10: IAS Role Descriptions and Tasks

Staff role	Role description	Tasks
IAS Administrators	This role performs day-to-day IAS administration tasks, such as controlling the IAS service and IAS configuration.	Start, stop, query, and configure the IAS service and make modifications to the IAS configuration database.
IAS Security Auditors	This role allows security auditors without administrative permissions to access security information.	Review RADIUS account and authentication log files for security events. When RADIUS authentication request logs are disabled, the IAS Security Auditors may need to review and save the System Event Log entries for IAS-related security events. This may require additional permissions.
IAS Backup Operators	This role allows Backup Operators to perform regular backups of the IAS servers. The backups include IAS configuration state and historical data.	Administer daily, weekly, and monthly backups of the IAS servers.
WLAN Helpdesk staff	Staff responsible for helping users to resolve issues related to WLAN access.	Review IAS events in the System Event Log related to user and device authentication or view the events as they are replicated to another system.

The following table details the resource permissions required to perform the various IAS server tasks.

Table 5.11: Permissions Required for IAS Server Tasks

Task	Group membership	Permission or rights required
Stop, start, query, and configure the IAS service	IAS Admins domain global group, which is a member of the local Administrators group on the IAS servers.	You can modify the service permissions in Windows Server 2003 using the SC command. Please consult with Microsoft support personnel before modifying the default permissions on the operating system components.
Modify the IAS configuration	IAS Admins domain global group, which is a member of the local Administrators group on IAS servers.	Permissions are required on the IAS database files located in the C:\WINDOWS\system32\ias directory, as well as to various registry keys under HKLM\System\CurrentControlSet\Services . By default, these permissions are granted to members of the Local/Builtin Administrators security group.

(continued)

<p>Access RADIUS request logs located on the IAS servers</p>	<p>IAS Security Auditors domain global group.</p>	<p>IAS Auditors must be able to read and delete RADIUS request log files located in the D:\IASLogs directory. Build guidance in this solution grants the NTFS Change permission to the IAS Security Auditors security group on this directory, and creates a share named IASLogs with Change share permission granted to the IAS Security Auditors security group.</p>
<p>Read and save IAS security events from the System Event Log</p>	<p>Local Administrators –or– Backup Operators on the IAS servers.</p>	<p>This solution provides guidance to enable RADIUS authentication logging to text files on disk. Therefore, IAS Auditors will typically not need to access IAS System Event Logs for RADIUS authentication security events. However, if you decide to disable RADIUS authentication logging, IAS Security Auditors must be able to read and save IAS events from the System Event Log. Archiving System Event Logs requires Administrator or Backup Operators membership.</p>
<p>Perform Daily, weekly, and monthly backups of the IAS servers</p>	<p>Backup Operators on the IAS servers.</p>	<p>Back up the IAS servers including the IAS configuration state and historical data such as RADIUS request logs. Membership in the Backup Operators security group enables access to the IAS database files located in the %systemroot%\system32\ias directory, various registry keys under HKLM\System\CurrentControlSet\Services, RADIUS request log files in D:\IASLogs, and IAS NETSH configuration text files located in D:\IASConfig.</p>
<p>Review IAS authentication events in the System Event Log for troubleshooting</p>	<p>Group membership with read permissions on the System Event Log.</p>	<p>Senior troubleshooting staff should be granted read permission on the Windows Server 2003 System Event Log to view and interpret IAS authentication reject events.</p>

Security Monitoring and Auditing

IAS is a component of your security infrastructure that you should proactively monitor. Security industry research has shown that successful attacks are typically preceded by a number of unsuccessful attacks. Proactive security monitoring of your IAS servers and their related logs for suspicious behavior is required to understand when your network is under attack.

The following table lists possible threats that you should monitor for in your IAS server infrastructure.

Table 5.12: IAS Server Infrastructure Threats

Threat/Vulnerability	Symptom	Monitoring tool
Authorization attempt using stolen credentials (such as those found on lost or stolen portable computer)	Authentication success/reject events (Source: IAS, ID 1 and 2) in the System Event Log or in the RADIUS authentication request logs indicating attempted use of certificates that have been revoked.	<ul style="list-style-type: none"> – MOM with a custom script written to parse Event Log entries for use of revoked certificates. – File parsing scripts or SQL Server tools that look for use of revoked certificates.
Attempted man-in-the-middle attack performed using a rogue wireless AP	System monitor counters on any IAS server showing excessive instances of the following: Bad Authenticators (Bad Message Authenticator attribute) or Invalid Requests (received from unknown RADIUS clients or servers).	MOM with a custom script to detect these system monitor counters and raise an alert.
Attempted DoS or buffer overflow against the IAS server service	System monitor counters on any IAS server showing excessive instances of the following: Malformed Packets (Packets containing malformed data), Unknown Type (non-RADIUS packets received), or Dropped Packets (packets dropped other than bad MAC/Malformed/Unknown).	MOM with a custom script to detect these system monitor counters and raise an alert.
Unauthorized authentication attempt	Repeated authentication failure events (Source: IAS, ID 2) in the System Event Log.	<ul style="list-style-type: none"> – MOM custom script to parse Event Log entries for patterns of excessive authentication rejections. – File parsing scripts or SQL Server tools to identify patterns of excessive authentication rejections.
Successful authentication using stolen credentials	RADIUS accounting logs indicate suspicious network activity.	<ul style="list-style-type: none"> – Microsoft Access to import logs and perform custom queries. – Reports to identify unusual network access information stored in a SQL Server database.

Beyond basic security monitoring, Microsoft recommends regularly auditing your IAS servers for potential security issues, and using your monitoring technology to clearly define and mitigate any vulnerabilities that you discover in your network infrastructure.

The following table details potential threats to an IAS server infrastructure and technologies that you can use to audit your IAS infrastructure for security issues.

Table 5.13: IAS Server Infrastructure Threats to Audit Proactively

Threat/Vulnerability	Symptom	Auditing tool
Weak permission to IAS configuration and historical data.	Unauthorized member of: the IAS Admins group, the IAS Security Auditors, or the Local Administrators group.	Active Directory and local security group auditing tools such as DumpSec by SomarSoft.
Attempts to hide unsuccessful authentication attempts.	System Event Log is unexpectedly cleared.	<ul style="list-style-type: none"> – Windows Event Log auditing using tools such as EventcombMT from the <i>Windows Server 2003 Resource Kit</i>. – Event log monitoring and alerting tool such as MOM.
Unauthorized modification to RADIUS account auditing and authentication logs.	Unexpected user ID shows write success in folder audit logs.	Windows file auditing and a monitoring tool such as MOM. To detect unauthorized file modification, you must enable file auditing.

Summary

This chapter explained the process of designing a RADIUS infrastructure to support a 802.1X-based secure wireless network. The RADIUS infrastructure design detailed in the chapter is flexible enough for you to extend it to a wide variety of future requirements. You also can use the infrastructure design for other types of network access management.

The design described in this chapter is used in subsequent chapters to implement the RADIUS infrastructure. The following chapter expands on the generic RADIUS design to describe the 802.1X settings and WLAN infrastructure required to implement the remaining components of the WLAN security infrastructure.

More Information

For more information about IAS, see the following resources:

- The [Windows Server 2003 Support Center](http://www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/entserver/default.asp) Web site at www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/entserver/default.asp.
The product documentation provides an overview of IAS features, basic instructions for configuration, and best practices for deployment.
- The [Microsoft Windows Server 2003 Technical Reference](http://www.microsoft.com/windows/reskits/default.asp) and the [Microsoft Windows Server 2003 Deployment Kit](http://www.microsoft.com/windows/reskits/default.asp) at www.microsoft.com/windows/reskits/default.asp.
- The “[IAS Technical Reference](http://www.microsoft.com/resources/documentation/windowsServ/2003/all/techref/en-us/W2K3TR_ias_intro.asp)” chapter of the *Microsoft Windows Server 2003 Technical Reference* at www.microsoft.com/resources/documentation/windowsServ/2003/all/techref/en-us/W2K3TR_ias_intro.asp.
This Technical Reference chapter provides technical information about IAS that is more comprehensive than the product documentation, and you can use it as a reference when more information is required.
- The “Deploying IAS” chapter of the *Deploying Network Services* guide of the [Microsoft Windows Server 2003 Deployment Kit](http://www.microsoft.com/windowsserver2003/techinfo/reskit/deploykit.mspx) at www.microsoft.com/windowsserver2003/techinfo/reskit/deploykit.mspx.
This deployment kit chapter contains deployment guidance for using IAS in a number of scenarios that are outside the scope of this secure wireless networking guidance, but that affect design decisions.

For more information about 802.1X WLAN technologies, see:

- The white paper “[Windows XP Wireless Deployment Technology and Component Overview](http://www.microsoft.com/technet/prodtechnol/winxpro/maintain/wificomp.mspx),” on Microsoft TechNet at www.microsoft.com/technet/prodtechnol/winxpro/maintain/wificomp.mspx.

6

Designing Wireless LAN Security Using 802.1X

Introduction

This chapter describes the architecture and design of the 802.1X-based secure wireless network components for the wireless LAN (WLAN) solution. The chapter presents the design decisions involved with securing wireless network components and the reasoning behind those decisions.

A important goal of the chapter is to help you determine the suitability of the design for your own organization. Where there are alternative design choices available, other relevant options are given alongside the option used in this solution. To help you understand the implications of the step without you having to refer to other documents, some topics are more detailed than might otherwise be necessary.

Chapter Prerequisites

Before reading this chapter, you should be familiar with both 802.11 wireless local area network (WLAN) concepts, 802.1X network access control, Remote Authentication Dial-In User Service (RADIUS) concepts, Microsoft® Internet Authentication Service (IAS), and WLAN deployment options using Microsoft Windows® XP Professional. You can familiarize yourself with these topics by reading the references listed at the end of this chapter in the “More Information” section. The *Microsoft Windows Server™ 2003 Resource Kit* and the *Microsoft Windows Server 2003 Deployment Kit* contain particularly valuable information.

Chapter Overview

The following flowchart illustrates the chapter structure.

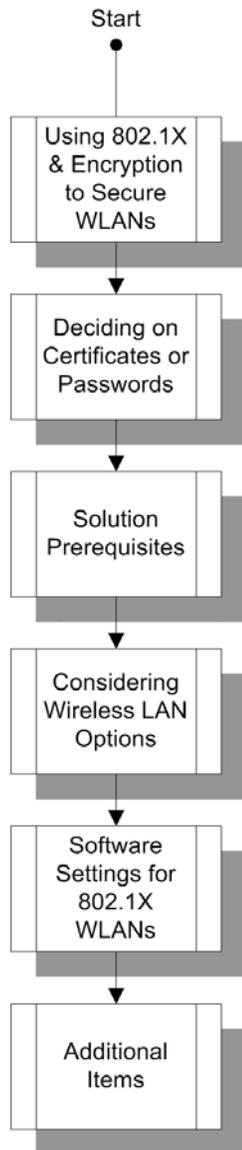


Figure 6.1
Planning WLAN security using 802.1X

This chapter covers six main steps:

1. **Using 802.1X and Encryption to Secure WLANs.** There are two main vulnerabilities to WLANs that anyone with a compatible WLAN adapter can exploit. This chapter explains both by telling you how to:
 - Enforce secure access to the network by configuring Institute of Electrical and Electronics Engineers (IEEE) 802.1X wireless access points (APs) as RADIUS clients to send access requests and accounting messages to RADIUS Servers. These RADIUS servers (running IAS) control network access through centralized remote access policies.
 - Protect data sent between the wireless devices and the wireless APs by using 128-bit Wired Equivalent Privacy (WEP) encryption or Wi-Fi Protected Access (WPA) encryption and integrity features built into 802.11X networking equipment. The data protection prevents interception and exploitation of radio-transmitted data.
2. **Deciding on Certificates or Passwords.** Microsoft offers native support for several types of authentication protocols that work with the 802.1X protocol. The most common forms of authentication credentials are passwords and digital certificates. The authentication method that you select for your organization can have a significant effect on the infrastructure required for your solution. This chapter helps you decide which method is best for your organization.
3. **Detailing Solution Prerequisites.** It is important to understand the solution prerequisites for the environment prior to starting your design. These include requirements for client computers, server infrastructure, and WLAN equipment. This section details these prerequisites.
4. **Considering WLAN Security Options.** Planning security options is complex and should include hardware purchasing representatives, security policy makers, usability representatives, network engineers, and network operations administrators. These experts should consider the following topics, which are detailed in this chapter:
 - Determining network authorization requirements
 - Choosing a client configuration strategy
 - Determining traffic encryption requirements
 - Producing a wireless network infrastructure design
 - Considerations related to wireless network Group Policy
5. **Determining Software Settings Required for 802.1X WLANs.** To achieve 802.1X WLAN security, you must configure both an IAS network access policy and an Active Directory® directory service Group Policy Object (GPO) for client computers. This section details how to achieve this.
6. **Considering Additional Factors.** This section briefly mentions topics to consider that are outside the scope of this solution but that may have an impact in your environment. These factors include:
 - Supporting roaming profiles and roving users
 - Supporting clients without wired LAN connections

Using 802.1X and Encryption to Secure WLANs

WLANs are becoming increasingly widespread with the adoption of industry standards such as IEEE 802.11 and 802.11b. WLANs allow users to roam around a building or campus and automatically connect to the network when they are in proximity to a wireless AP.

While providing convenience, WLANs pose the following security risks:

- Anyone who has a compatible WLAN adapter can gain access to the network.
- Wireless networking signals use radio waves to send and receive information. Anyone within an appropriate distance to a wireless AP can detect and receive all data sent to and from the wireless AP.

To counter the first security risk, you can configure IEEE 802.1X wireless APs as RADIUS clients to send access requests and accounting messages to RADIUS servers running IAS. IAS performs authentication of users and devices, and controls network access through centralized remote access policies.

To counter the second security risk, you can protect data sent between the wireless devices and the wireless APs by using 128-bit WEP encryption or WPA encryption capabilities built into 802.11 networking equipment.

Static WEP has serious design flaws and includes no native encryption key management to allow the keys to be regularly updated. This can expose encryption keys to determined malicious users. IAS enables strong WEP keys to be dynamically assigned to client computers running Windows XP during certificate-based authentication. In addition, WEP keys can be regenerated at regular intervals to thwart attack tools designed to discover those keys.

WPA is a subset of the forthcoming 802.11i security standard for 802.11-based WLAN equipment. WPA includes enhanced encryption designed to address security issues with static WEP. The solution presented in this guide also is suitable for use with WPA (this requires WPA capable hardware and an update to the Windows XP client).

Deciding on Certificates or Passwords

Microsoft offers native support for several different authentication methods that you can use with the 802.1X protocol. Most commonly, organizations select WLAN client authentication methods based on either passwords or certificate-based credentials.

As stated earlier, the authentication method selected can have a significant effect on the infrastructure required for your solution. The 802.1X standard uses an authentication scheme called Extensible Authentication Protocol (EAP) that allows you to “plug in” different authentication types.

The following table illustrates the EAP types that you can use in a Microsoft 802.1X infrastructure, and the advantages and disadvantages of each type.

Table 6.1: Advantages and Disadvantages of EAP Types

Feature	PEAP	EAP-TLS	EAP-MD5
Mutual authentication	Mutual authentication.	Mutual authentication.	Only client authentication.
Dynamic key generation and scheduled regeneration.	Generated during authentication and regenerated at timed intervals.	Generated during authentication and regenerated at timed intervals.	No dynamic key generation or regeneration: Relies on static keys.
Security technology level	Can use strong password authentication or digital certificates.	Strongest authentication.	Weak security technology.
User credential protection	Protected by Transport Layer Security (TLS) tunnel.	Certificate-based authentication protected by Transport Layer Security (TLS) tunnel.	Open to dictionary attack.
Ease of implementation	Widely supported and offered natively in Windows clients.	Requires a Public Key Infrastructure (PKI). Widely supported and offered natively in Windows clients.	Simple, but not recommended for wireless.
Credentials flexibility	Any approved EAP with TLS tunnel, including EAP–MSCHAPv2 (passwords-base method).	Only digital certificates.	Password only.

The recommended EAP type for performing certificate-based client authentication is EAP–TLS, and the recommended EAP type for performing password-based client authentication is EAP–MSCHAPv2 within Protected Extensible Authentication Protocol (PEAP), also known as PEAP–EAP–MSCHAPv2.

Password-based 802.1X authentication using PEAP and MSCHAPv2 is a low cost and robust solution. It is suitable for organizations that do not currently have a certificate infrastructure in place and do not need certificates for other purposes, such as using the Encrypting File System (EFS) or a virtual private network (VPN). You can easily migrate from password-based 802.1X authentication to certificate-based authentication. This provides flexibility for your organization to later change from one authentication method to the other.

Note that even with a password solution using PEAP, a certificate is required on each RADIUS server. You should weigh the costs involved with purchasing server certificates from a commercial certificate provider against the value that a certificate infrastructure will bring to your organization.

This solution uses certificate-based client authentication because of the greater security level that the authentication method provides. The authentication method uses the Extensible Authentication Protocol–Transport Layer Security (EAP–TLS) protocol.

For guidance on how to deploy a password-based 802.1X solution that uses PEAP and MSCHAPv2, see the discussion in Chapter 2 "Deciding On a Secure Wireless Networking Strategy," and the reference at the end of this chapter to the companion solution guide, *Securing Wireless LANs with PEAP and Passwords*.

Solution Prerequisites

It is important that you understand the prerequisites for the environment for this solution before starting the design. This section details these prerequisites.

Client Computer Requirements

This solution has been designed and tested using Windows XP Professional with Service Pack (SP) 1. Windows XP with SP1 provides certain 802.1X and WLAN feature functionality that is required to achieve an extremely low cost and easily managed solution.

Testing this solution included client computers running both Windows XP Professional and Windows XP, Tablet PC Edition. Both Windows XP editions provide automatic certificate enrollment and renewal of computer and user WLAN authentication certificates that are required for EAP-TLS client authentication. This feature alone significantly reduces the cost typically associated with certificates, and thus a certificate-based 802.1X solution.

Microsoft also provides 802.1X clients for Windows 2000 (available as a free download), and Windows 9x and Microsoft Windows NT® 4.0 (available free of charge to customers with a support agreement). However, these client types were not tested in this version of the solution.

Required Server Infrastructure

This solution depends on the Certificate Services and IAS components of Windows Server 2003. There are features of Certificate Services and IAS that have been designed explicitly for 802.1X-based WLANs. Some of the features used in this solution include editable certificate templates and remote access policy settings that allow you to simplify deployment settings required for the 802.1X protocol.

The solution was designed with both Windows Server 2003 and Windows 2000 Active Directory environments in mind. However, the solution was tested only with Windows Server 2003 domain controllers. If you choose to, you can install IAS on to existing domain controllers. For a detailed discussion of co-location considerations related to this option, see Chapter 5, “Designing a RADIUS Infrastructure for Wireless LAN Security.”

Required WLAN Equipment

This solution assumes that your organization has already deployed a well designed and fully functional WLAN infrastructure. This guide does not include any instructions on wireless network design, such as wireless AP positioning and channel selection. If your organization has not deployed a WLAN infrastructure, ensure that you have the appropriate level of expertise available to accomplish this before you start deploying the WLAN security components.

The network hardware must support 802.1X, and 128-bit WEP for encryption. In this solution guidance, it is assumed that the WLAN infrastructure is operating without error, and that either no security controls are enabled, or only basic 802.11 security controls are enabled. Migrating from either a Shared Key (static WEP) WLAN or an Open System (unsecured) 802.11 WLAN to this solution is very similar. You should be able to accomplish either type of migration without any significant problems.

Considering WLAN Security Options

If you have not already, you should spend time planning a WLAN security policy for your organization. Your planning discussions should include representatives from hardware purchasing, security policy, usability, network engineering, and network operations. Your security policy discussions with these people should include how your organization will address the threats it faces and what security controls you will use to mitigate them.

It is also important to document your WLAN security policy and make it available and visible to all of your network users. This solution provides security controls to mitigate the risk associated with modern WLAN technology. However, the solution cannot mitigate the risks of users in your organization who are performing unsecured, improvised networking, and deploying rogue wireless APs.

Selecting User- and Computer-Based Authentication

User authentication is a natural choice when considering identification to WLAN infrastructure. However, in most cases you also will want to implement computer (or device) authentication to ensure a comprehensive security solution for your WLAN.

There are a number of features in Windows XP Professional that will only work correctly with an active network connection. Using 802.1X computer authentication ensures that the WLAN network connection is established during the computer startup sequence before users see the initial Windows logon screen. The computer re-authenticates to the WLAN after the user logs off, ensuring that there is always a connection to the network.

Table 6.2: Reasons for using Computer Authentication

Feature	Scenario requiring computer authentication
Active Directory computer Group Policy	Computer-based Group Policy is applied during computer start up and at timed intervals—even when no one is logged in to the Windows operating system.
Network logon scripts	Network logon scripts are run during initial user logon.
Systems management agents	Systems management agents, such as those that come with Microsoft Systems Management Server (SMS), frequently require network access without user intervention.
Remote Desktop Connection	Computers are accessible via the Windows Remote Desktop Connection when no one is logged on to the computer.
Shared folders	Files and folders shared from a computer are still available, even when no user is logged on to it.

The best strategy is to use user-based authentication when possible, and computer-based authentication when required. This solution uses the default behavior of the Windows XP Professional 802.1X client. This is to perform computer authentication when no user is logged on at the computer console, user authentication when a user logs on to Windows, and computer authentication again when the user log off. This ensures that authentication to the network uses user account credentials where possible for accountability, while still it ensures that Windows features that require network access continue to work properly when no one is logged on.

Validating Certificate Based Credentials

It is important to check for valid credentials as part of your certificate-based WLAN authentication strategy. Checking for revoked certificates allows you to block the use of client certificates stored on lost or stolen computer equipment. Forcing clients to verify server certificates helps prevent sophisticated man-in-the-middle attacks involving rogue APs and RADIUS servers.

Windows provides extensive support for validating certificates when performing certificate-based operations. Both IAS and the Windows XP Professional 802.1X features use this support to ensure that the certificates used for EAP-TLS are valid and represent a trusted security principal.

Table 6.3: IAS Validation of Client Certificate Credentials

IAS validation of client certificate credentials	Default behavior	Settings used in this solution
Check that the certificate is within its validity dates.	Enabled	No change
Check that it is possible to build a chain from the certificate to a trusted root.	Enabled	No change
Check that the required key usages and application policies are present in the certificate.	Enabled	No change
Check that the client proves ownership by signing with a private key.	Enabled	No change
Check that the certificate is not revoked.	Enabled	No change

Windows XP Professional also performs the following IAS server credentials validation by default.

Table 6.4: Windows XP Validation of IAS Certificate Credentials

Windows XP validation of server certificate credentials	Default behavior	Settings used in this solution
Check that the certificate is within its validity dates.	Enabled	No change
Check that it is possible to build a chain from the certificate to a trusted root.	Enabled	No change
Check that the required key usages and application policies are present in the certificate.	Enabled	No change
Check that the server proves ownership by signing with a private key.	Enabled	No change

When authenticating to the WLAN, the client computer cannot perform a full revocation check of the server certificate, because network access is not available prior to the completion of a successful authentication.

You should also consider enabling the following additional credential validation options (that the client performs) to increase the security of the validity check.

Table 6.5: Advanced Windows XP Validation of IAS Certificate Credentials

Windows XP validation of server certificate credentials	Default behavior	Settings used in this solution
Subject of the certificate matches a Domain Name System (DNS) string value that you can configure on the client.	Not enabled	No change
Explicit selection of trusted root CAs that the server certificate may chain.	Not enabled	Enabled

Note that the subject name checking option on client computers will produce a trust decision prompt to users. In addition, you need to implement a management process to keep the permitted server certificate subject names up to date on the WLAN clients. You can do this using the Wireless Network Policies GPO settings. For these reasons, this solution does not implement this option. If you operate a high security environment, you may want to consider the threats that rogue IAS servers with wireless APs pose when deciding whether you need this extra level of validation.

Explicit selection of trusted root authorities minimizes the risk of forged server certificates from alternate CAs in the trusted root store. However, this setting requires an additional management process to ensure that changes to trusted root authority certificates are reflected in the WLAN client settings. These settings are also deployed using Wireless Network Policies GPO settings.

Determining Network Authorization Requirements

The key goals to achieve when designing network access management policies are to match organizational security policies as closely as possible and minimize management cost. A centralized representation of network authorization policy such as IAS remote access policies is well suited for this task.

Note: This solution uses network authorization administration through IAS remote access policy. For more information about the decisions involved to select a remote access policy administration model, refer to the resources listed in the “More Information” section at the end of this chapter.

Flexible yet simple management of network authorization should be a primary goal for any organization. Minimizing the number of remote access policies while still ensuring that all organizational security policies are represented is the key to achieving simplified management.

Determining Connection Criteria

IAS remote access policies can either allow or deny connections. A policy contains a set of criteria against which to match each connection attempt. The first policy found that matches a particular connection will be used to allow or deny access. The main connection attributes against which a policy can be matched are the following:

- Group membership
- Type of connection
- Time of day
- Authentication methods

In addition, you can specify many other advanced filter criteria such as the following:

- Access server identity
- Access client phone number or Media Access Control (MAC) address
- Whether to ignore user account dial-in properties
- Whether to allow unauthenticated access

This solution uses IAS connection criteria based on domain security group and source network type (rather than time of day, authentication method, or other condition). This ensures that the remote access policies are created for a particular type of network access (such as, wireless, VPN, wired, or dial-up) and client grouping, while keeping the policy criteria broad enough to minimize the number of policies required.

Note: This solution uses custom security groups (Remote Access Policy–Wireless Users, and Remote Access Policy–Wireless Computers) to restrict which users and computers are allowed access to the WLAN. If you want all of your domain users and computers to be able to access the WLAN, you can add the Domain Users and Domain Computers groups to these custom security groups to simplify administration.

Determining Connection Restrictions

Once a connection is authorized, the remote access policy also specifies connection restrictions and other attributes to be applied to that connection. These include the following:

- Idle time-out
- Maximum session time
- Encryption strength
- IP packet filters

In addition, you can apply the following attributes to the connection:

- IP address for Point-to-Point Protocol (PPP) connections
- Static routes

One major determining factor for the number of remote access policies that your organization will require is the number of different user types who need access to the wireless network. For example, full-time staff in many organizations requires full, unrestricted access to the entire corporate network. However, contractors and business partners may only require access to specific applications on specific network subnets.

Connection restriction profiles are unique to each remote access policy. Therefore, if multiple types of connection restriction profiles are required, so are multiple remote access policies.

The following table illustrates examples of various types of users and the some of the connection restrictions that you can apply through remote access policy.

Table 6.6: Examples of WLAN Connection Restrictions

Type of user group	Connection restriction example
Full-time staff	Authenticated, unrestricted access to the corporate network.
Contractors and business partners	Authenticated, restricted access to specific networks and applications.
Visiting guests	Unauthenticated access to Internet-only segments for Web browsing or VPN access to source organization.

This solution only configures support for authenticated, full-time staff. Thus, only a single remote access policy with a simplified connection restriction profile is required. If your organization has additional requirements, such as a requirement to support restricted access users to the wireless network, you must add remote access policies to provide for this. See the references in the “More Information” section at the end of this chapter for more information about planning additional types of remote access policies.

Choosing a Client Configuration Strategy

Automating client computer configuration settings is an essential step to reducing the cost of deploying wireless networking security, and minimizing support issues that result from incorrectly configured settings.

Windows XP Professional has sophisticated features for reducing the need for manual configuration and reconfiguration of wireless network and 802.1X security settings on the WLAN clients. Windows Server 2003 adds the capability to completely automate client configuration using the Wireless Networking Policies settings in Group Policy. This solution uses the Wireless Networking Policy feature in Windows Server 2003 to automatically configure all wireless network clients.

You can deploy these GPO settings to client computers even before distributing WLAN network interface cards (NIC). This allows end users to simply install the wireless NICs and, using the WLAN settings deployed via the GPO, automatically connect to the secure 802.1X WLAN.

Determining Traffic Encryption Requirements

You should keep up to date on the evolving threats to 802.11 WEP encryption when determining your strategy for protecting your WLAN traffic. As discussed earlier, it is possible for determined malicious users to take advantage of cryptographic flaws in WEP to perform attacks that can disclose the WEP encryption key. To perform this kind of attack an intruder needs to capture several million packets that have been secured with the same encryption key.

The best strategy for mitigating threats to the security of a WEP-based WLAN is to ensure that keys used to encrypt network traffic are refreshed periodically. This can be done at a frequency that prevents an attacker ever being able to capture enough traffic to

perform a successful attack. You can accomplish this by setting the IAS RADIUS options to enforce automatic client re-authentication, which initiates WEP key regeneration.

This solution configures IAS RADIUS options to enforce Windows XP client re-authentication every 10 minutes to ensure short-lived WEP session keys. This decision was based on known WEP attack tools and scenarios at the time of writing. Basic WEP encryption includes flaws such as poor initialization vector (IV) sequencing that an attacker can exploit to more quickly compromise keys. Ensure that your APs have the capability to generate less predictable and therefore stronger IVs.

Important: Using a session timeout of 10 minutes may be too short for many scenarios. Short timeouts place a high load on the IAS servers. Short timeouts also increase the possibility of making your IAS server temporarily unavailable, which will result in disconnecting clients from the WLAN. For these reasons, you can use a longer timeout of 60 minutes without significantly compromising the WLAN security.

The use of EAP–TLS ensures that unique WEP session keys are generated for each client during the TLS negotiation process, and that they are transported safely across the network between solution components. Unlike static WEP, the encryption keys are never reused and never shared between clients.

Choosing a WLAN Migration Strategy

If you have an existing wireless network in place, you should plan a migration strategy up front to ensure minimal disruption to users and the environment.

Studies have shown that many organizations have 802.11-based WLANs in place and are operating without network authentication or data protection. This type of 802.11 network security strategy is called *Open System* authentication. Other organizations have implemented static key network authentication and encryption. This type of 802.11 network security is also known as *Shared Key* authentication.

Migrations from Open System or Shared Key network security models to 802.1X security are very similar. The primary difference is that Shared Key security provides some security protection, thus migration schedules from Shared Key security may be more relaxed for some organizations.

Migration from either of these authentication strategies to the 802.1X security model typically involves the following steps:

1. **Deployment of certificates to computers and users**— This should be done in advance of the 802.1X deployment to ensure that the certificates are deployed to mobile computers that only occasionally connect to the LAN.
2. **Configuration of wireless network remote access policies on IAS servers** — This involves guidance on configuring a wireless remote access policy.

3. **Deployment of wireless network configurations on client computers**—The new 802.1X enabled network will typically require a new network Service Set Identifier (SSID). You can deploy network settings for this SSID using Active Directory Group Policy. WLAN group policy must be deployed sufficiently in advance of wireless AP reconfiguration to ensure that mobile computers that connect to the LAN infrequently will have received the WLAN GPO settings.
4. **Configuration of wireless APs to require 802.1X security**—This step is typically performed on a location-by-location basis, such as by each building in the environment or campus. You should plan appropriate roll-back procedures in the event of unexpected behavior, and staff Help desks appropriately to handle associated support calls.

As with all migration strategies, careful planning is essential. Configuring client computers and wireless APs erroneously can cause disruptive changes to your environment. You should test the intended changes thoroughly prior to deployment.

Note: Some wireless APs support configuring one 802.11 radio for static WEP security and another 802.11 radio for 802.1X. However, 802.11 channel separation issues may make this strategy impractical for many organizations.

Obtain commitment from your WLAN equipment vendors to support wireless AP and wireless NIC upgrades to WPA. Microsoft has released an update to Windows XP that provides support for WPA (it is included in SP2). Also, plan to upgrade your WLAN infrastructure in the future to support 802.11i, which will likely require you to update the firmware on your APs and wireless NICs.

This guidance does not include detailed migration planning for production WLANs that use proprietary security or Open System/Shared Key security. For assistance with detailed migration planning from production WLANs, see your Microsoft partner or contact your Microsoft Account Executive who can connect you with the appropriate partner or Microsoft Consulting Services professionals.

Wireless Network Infrastructure Design

General discussion on 802.11-based WLAN equipment and network design is beyond the scope of this guidance. The WLAN chapter in the *Microsoft Windows Server 2003 Deployment Kit* provides general guidance on this topic.

Effort has been made to ensure that this solution will work on a broad variety of products from various network equipment vendors. Configuration planning and procedures for specific wireless AP products is beyond the scope of this solution.

However, when you decide on the security settings and security management of your 802.11 equipment, educate yourself on the following topics by using documentation available from your hardware manufacturer:

- **SSID name and default password**—This topic includes changing the default SSID on all APs, and choosing a strategy regarding whether to configure APs to broadcast their SSIDs.
- **Changing default console password and Simple Network Management Protocol (SNMP) strings**—This topic includes changing the default management passwords and access strings on wireless APs and maintaining them over time.
- **Secure administration of wireless APs**—This topic includes using secure communications when performing administration on wireless APs by using protocols such as Secure Shell (SSH) or Hypertext Transfer Protocol (HTTP) with SSL or TLS.

- **RADIUS client settings**—This topic includes configuring your APs to use RADIUS servers for authentication and accounting. Discussion on configuring APs for a primary and secondary RADIUS Server, use of strong RADIUS secrets, and the Message Authenticator attribute appears in Chapter 5, “Designing a RADIUS Infrastructure for Wireless LAN Security,” and Chapter 9, “Implementing the Wireless LAN Security Infrastructure.” Chapter 9 includes instructions on using a supplied script to generate strong RADIUS secrets.
- **Virtual local area network (VLAN) switching and traffic filtering**—This discusses using network VLANs to restrict access to various types of users in different scenarios. IAS can provide RADIUS values based on remote access policy to assist with automated selection of appropriate VLANs during client connection. For more information about the IAS options for specifying VLANs to wireless APs, consult the references listed in the More Information section at the end of this chapter.
- **Tactics for limiting leakage of wireless LAN radio transmissions outside building boundaries**—This topic includes such items as avoiding placement of APs against exterior walls or windows. It also includes reducing the broadcast strength of the APs when possible to keep coverage within the necessary area, and avoid coverage of unintended areas, such as parking lots.
- **Rogue wireless AP detection**—This topic includes actively and regularly scanning for rogue APs on the corporate network using available Windows XP and Windows CE-based WLAN management tools, such as NetStumbler or AirMagnet.

For assistance with these topics and wireless AP configuration, consult your vendor documentation or enlist the assistance of an equipment specialist. Some of these items are discussed in the "802.11 Wireless Network Technical Reference" topic in the *Windows Server 2003 Technical Reference* in the "More Information" section at the end of this chapter.

Wireless Network Group Policy Considerations

Consult with your domain GPO administrators to determine the strategy for applying Wireless Network Group Policy to client computers. The following table lists the key items that you need to decide on for your own environment and the settings that have been selected in this solution.

Table 6.7: Wireless Network Policy Planning

Wireless network policy consideration	Solution strategy	Solution details
Policy application criteria	Active Directory security group filtering to include selected computers.	Wireless Network Policy—Computer global group.
Number of GPOs required	Single Wireless Network Policy.	The GPO used in this solution is called "Wireless Network Policy."
GPO location	Created and applied from the domain object.	woodgrovebank.com.
Number of WLAN profiles configured within the policy	One WLAN profile is configured for organizations implementing 802.1X.	The GPO contains one WLAN profile for environments that have not entered their WLAN into production use. You may add additional WLAN profiles as required to suit a phased migration from a legacy production WLAN.

Note: This solution grants a custom security group (Wireless Network Policy—Computer) "Apply Policy" permission on the "Wireless Network Policy" GPO. Membership of this group therefore determines which computers receive the WLAN GPO settings. If you want to allow all computers to receive the WLAN configuration settings, you can add the Domain Computers or Authenticated Users group to this group to simplify administration. However, you should be aware that doing this will apply the policy settings to all servers and clients in the domain (Domain Computers) or forest (Authenticated Users).

This solution uses a simple Wireless Network policy management strategy by creating a single GPO linked to the domain object. This means that any computer in the domain who is also a member of the "Wireless Network Policy—Computer" group will receive the policy settings. You may want to consider more sophisticated Group Policy management standards and apply your Wireless Network policies accordingly. Most organizations, however, will need only one Wireless Network Policy GPO (although you may choose to link it to a location other than the domain object).

Determining Software Settings Required for 802.1X WLANs

There are two major software solution components that you must configure to achieve 802.1X WLAN security:

- IAS network access policy
- Active Directory Group Policy for client computers

IAS network access policy is at the heart of your network access management solution. Settings that represent your WLAN security policy are deployed on each IAS-based RADIUS server in an automated fashion. This policy includes settings for:

- Remote access policies
- Connection request policies

Remote access policies enforce access to your networks via your wireless APs. Connection request policies determine the handling of RADIUS requests from various wireless APs configured as RADIUS clients.

Active Directory Group Policy for client computers contains all of the settings that will be deployed to Windows XP-based client computers. This Group Policy affects client interaction with the wireless APs, the RADIUS server, and other wireless networks.

Configuring Remote Access Policies

You must plan the creation of remote access policies on IAS servers to satisfy your organizational wireless network access strategy. Creating and configuring remote access policies involves establishing the following three setting types for each policy:

- Policy conditions
- Policy permission
- Policy profile

This solution uses a single remote access policy that grants unrestricted access to the wireless network for full-time employees. The following table details the remote access policy conditions for this solution.

Table 6.8: Remote Access Policy Conditions

Policy condition	Matching condition	comment
NAS-Port-Type	Wireless-Other or Wireless-IEEE 802.11	This identifies incoming requests as originating from wireless AP hardware.
Windows-Group	Membership in Remote Access Policy—Wireless Access security group	This is a domain universal security group that contains nested global groups for users and computers that will receive access to the WLAN.

The policy permission for the remote access policy in this solution is set to **Grant**, and user accounts are configured with the **Control access through Remote Access Policy** setting.

The following table details the remote access policy profile options that this solution uses. You may need to add additional settings to suit your particular environment.

Table 6.9: Remote Access Policy Profile Options

Profile option	Profile setting	Comment
Dial-in Constraints—Defines minutes that clients can be connected (Session-Timeout)	10 minutes	This setting forces client computers to re-authenticate and create new encryption keys every 10 minutes.
Authentication—EAP Methods	Smart card or other certificate	This setting selects EAP—TLS as the EAP type for the wireless profile.
Ignore—User—Dial-in-Properties RADIUS attribute	Attribute set to True	This attribute ensures that dial-in settings (such as callback) on Active Directory user accounts are not sent to wireless APs. This is done to avoid issues with some network access products.
Termination-Action RADIUS attribute	Attribute set to RADIUS Request	This attribute ensures that when clients are forced to re-authenticate, wireless APs do not disconnect them.

Important: Using a session timeout of 10 minutes may be too short for many scenarios. Short timeouts place a high load on the IAS servers. Short timeouts also increase the possibility of making your IAS server temporarily unavailable, which will result in disconnecting clients from the WLAN much sooner. For these reasons, you can use a longer timeout of 60 minutes without significantly compromising the WLAN security.

Configuring Connection Request Policies

You should plan how IAS will handle RADIUS requests when it is operating as a RADIUS server and as a RADIUS Proxy. Discussion of selecting the RADIUS role for IAS appears in Chapter 5, “Designing a RADIUS Infrastructure for Wireless LAN Security.”

Regardless of which role you select for the IAS server, you must configure the following components of the connection request policies before you can achieve wireless network access:

- Policy conditions
- Policy profile

This solution uses IAS as a RADIUS server, thus connection requests are processed locally on each server. The settings in the following table illustrate policy conditions configured in the connection request policy for this solution.

Note: The connection request policy settings in this solution use the defaults that install with IAS in Windows Server 2003.

Table 6.10: Connection Request Policy Conditions

Policy condition	Matching condition	Comment
Date-And-Time-Restrictions	“Sun 00:00–24:00; Mon 00:00–24:00; Tue 00:00–24:00; Wed 00:00–24:00; Thu 00:00–24:00; Fri 00:00–24:00; Sat 00:00–24:00”	This condition of the default connection request policy ensures that connection requests arriving at any time match the policy.

The following table illustrates profile settings used in the connection request policy for this solution.

Table 6.11: Connection Request Policy Profile Settings

Profile option	Profile setting	Comment
Authentication	Authenticate requests on this server	This setting ensures that requests are authenticated directly against Active Directory instead of being forwarded to additional RADIUS servers.

Configuring Group Policy for Client Computers

Planning is required prior to using Active Directory Group Policy to configure WLAN settings and 802.1X security settings on client computers via Wireless Network (IEEE 802.11) Policies. This section details a number of the settings in this solution and the reasons to include them. You will find the Wireless Network (IEEE 802.11) Policies on the \Computer configuration\Windows Settings\Security Settings\Wireless Network (IEEE 802.11) Policies object within the Group Policy Object Editor.

Configuring Wireless Network Settings

You configure WLAN settings by editing the properties of the WLAN policy objects within Group Policy. These include the following setting types:

- General settings
- Preferred networks
- Network properties

The following table illustrates general settings in the Wireless Networks Policy for this solution.

Table 6.12: Wireless Network Policies General Settings

Option	Setting	Comment
Name	Client Computer Wireless Configuration	You can change this value to match your organizational naming standards.
Network to access	Any available network (access point preferred)	This setting prevents client computers connecting to other computers that are configured with the same SSID as your 802.1X enabled network. However, allowing ad-hoc networking enables employees to use other networks when required (such as at home), and therefore this setting has been left enabled.
Automatically connect to non-preferred networks	Cleared	Windows XP Professional automatically provides users with notification of available wireless networks without automatically connecting to them. Notifying users without automatically connecting them strikes a balance between security and usability.

You must create an entry for your 802.1X WLAN SSID on the **Preferred Networks** tab in the Wireless Networks Policy. Once a preferred network is created, you must edit the network properties from the defaults.

The following table details network properties settings for the newly enabled 802.1X network in the Wireless Networks Policy for this solution.

Table 6.13: Wireless Network Policy Properties Settings

Option	Setting	Comment
Name	MSSWLAN	Change this value to match your organization's naming standards. However, be sure to choose a name that is different from any existing production WLAN.
Wireless network key (WEP)– Data encryption (WEP enabled)	Selected	Encryption is essential to protect the privacy of wireless network traffic on 802.11 networks. If you support 802.11 networks, ensure that they are protected by WEP or some other form of encryption.
Wireless network key (WEP)– Network authentication (Shared mode)	Cleared	Shared key 802.11 wireless is a security strategy based on static WEP keys. This solution uses 802.1X to provide RADIUS authentication against Active Directory, and, therefore, this option is cleared.
Wireless network key (WEP)– The network key is provided automatically	Selected	Having this setting enabled allows 802.1X to automatically provide dynamic WEP session keys for network traffic encryption.
This is a computer-to-computer (ad-hoc) network; wireless access points are not used	Cleared	This setting in the solution uses 802.11 WLAN infrastructure mode with wireless APs configured for 802.1X, not point-to-point ad-hoc networking.

Configuring 802.1X Settings on Client Computers

You configure 802.1X settings via your Wireless Network Policy, which includes the following setting types:

- 802.1X parameters
- EAP type
- Credentials validation
- Computer authentication behavior

The following table details 802.1X settings for the newly enabled 802.1X network in the Wireless Networks Policy for this solution.

Table 6.14: Wireless Network Policy 802.1X Settings

Option	Setting	Comment
Enable network access control using IEEE 802.1X	Enabled	This option enables the client computer to participate in networks secured using 802.1X.
EAPOL-Start message	Transmit	This message tells the client to start the authentication process.
Parameters (seconds)–Max Start	3	This value determines the number of successive EAP over LAN (EAPOL)–Start messages that the client will transmit after not receiving a response. This value should not be changed from the default value unless required.
Parameters (seconds)–Held Period	60	This value determines the amount of time the client will wait before re-attempting a failed 802.1X authentication. This value should not be changed from the default value unless required.
Parameters (seconds)–Start period	60	This value determines the time interval for resending EAPOL-Start messages. This value should not be changed from the default value unless required.
Authentication period	30	This value determines the time interval for resending 802.1X request messages after not receiving a response. This value should not be changed from the default value unless required.
EAP type	Smart card or other certificate	This option specifies EAP–TLS as the EAP type.

The following details the EAP settings for the newly enabled 802.1X network in the Wireless Networks Policy for this solution.

Table 6.15: Wireless Network Policies EAP Settings

Option	Setting	Comment
When connecting	Use a certificate on this computer	This option specifies the use of software-based certificates and private keys rather than smart card-based credentials.
Use a certificate on this computer—Use simple certificate selection (Recommended)	Selected	This enabled option determines that Windows will attempt to choose the correct certificate based on certificate properties. You can turn this option off to enable manual selection of the correct certificate when troubleshooting.
Validate server certificate	Selected	This enabled option determines whether the certificate presented to the client during EAP—TLS authentication is valid (that the IAS server's correct DNS name is in the certificate, and that the certificate for server authentication has not expired and chains to a root CA in the client's certificate store).
Validate server certificate—Connect to these servers	Cleared	When enabled, this option allows checking of the fully qualified domain name (FQDN) suffix in the subject field of the server certificate. Enabling this option will produce a text balloon box on client computers prompting for trust approval of the IAS server. You should evaluate usability versus security when selecting this option.
Validate server certificate—Connect to these servers—Value	Blank	This value specifies the FQDN suffix that must match the subject information in the certificate presented to the client during EAP—TLS authentication.
Trusted root certification authorities (CAs)	CompanyCA selected	This option allows administrators to specify the trusted root CAs to which 802.1X server certificate credentials are allowed to chain. You should select your own Trusted root CAs for this option.
Use a different name for the connection	Cleared	When enabled, this option allows users to specify a user name other than that contained within the certificate presented during EAP—TLS authentication. This option is disabled for this solution.

The following table details computer authentication settings for the newly enabled 802.1X network in the Wireless Networks Policy for the solution.

Table 6.16: 802.1X Computer Authentication Behavior Options

Option	Setting	Comment
Authenticate as guest when user or computer information is unavailable	Cleared	When enabled, this setting determines whether the computer will attempt to authenticate as guest when no credentials are available. This is most useful for public WLAN scenarios or visitors to your organization.
Authenticate as computer when computer information is available	Selected	It is essential to enable this setting to ensure that computer authentication occurs when users are not interactively logged on to the computer.
Computer authentication	With user re-authentication	This default option ensures that user credentials are used whenever possible. However, when users are not logged on to the computer, computer credentials are used to ensure that a network connection is available at all times.

Additional Considerations

This section briefly mentions additional topics to consider that are outside the scope of this solution but that may have an impact in your environment.

Supporting Roaming Profiles and Roving Users

EAP–TLS-based 802.1X components of Windows depend on certificates and private key information being available within the certificate store on your client computers. For most organizations, wireless users have portable computers or Tablet PCs that travel with them, and thus the certificates and key information is always available.

However, if your WLAN strategy includes users who share computers, you will want to consider implementing roaming profiles. You can use roaming profiles to ensure that private key information and certificates are always available for use within your 802.1X-based environment.

Alternatively, Wireless Network Policy allows you to configure computers running Windows XP to perform computer-only authentication. Although not implemented in this solution, this may be useful for some organizations.

Supporting Clients Without Wired LAN Connections

Although not common for large organizations, some environments may not have a wired LAN at all. Wired LAN infrastructure is required to join client computers to a domain, and receive certificates and Group Policy. Without computer and user certificates, and appropriate client WLAN configuration, users cannot access the 802.1X-based WLAN. This issue also arises when an organization has deployed a wired network that always requires 802.1X authentication.

If you have such an environment, you should consider a strategy whereby users provide a password-based credential using PEAP-MSCHAPv2 to the RADIUS server for connection to a controlled VLAN with the Certificate Services Web enrollment pages. From there, users can enroll and install certificates to gain full access to the corporate WLAN using EAP–TLS. Such a strategy is currently outside the scope of this solution and would require an additional remote access policy on the IAS servers, in addition to specific VLAN design.

Summary

This chapter described the process of designing WLAN security using the 802.1X protocol, which included determining WLAN prerequisites, considering security options, establishing strategy, and additional considerations. Once you have established your design based on the options discussed in this chapter, you are in a position to deploy 802.1X-based WLAN security in your environment.

The design in this chapter will be used in the later Build and Operations chapters for this solution to implement the 802.1X WLAN security infrastructure.

More Information

For more information about secure wireless networking, see the following resources:

- The [Product Documentation for Windows Server 2003](http://www.microsoft.com/windowsserver2003/proddoc/default.mspx) Web site at www.microsoft.com/windowsserver2003/proddoc/default.mspx.
The product documentation provides an overview of IAS features, basic instructions for configuration, and best practices for deployment.
- The Microsoft solution for [Securing Wireless LANs with PEAP and Passwords](http://go.microsoft.com/fwlink/?LinkId=23459) is available at <http://go.microsoft.com/fwlink/?LinkId=23459>.
- The “[IAS Technical Reference](#)” chapter of the *Microsoft Windows Server 2003 Technical Reference* at www.microsoft.com/resources/documentation/windowsserv/2003/all/techref/en-us/W2K3TR_ias_intro.asp.
This resource kit chapter provides technical information about IAS that is more detailed than the product documentation, and you can use it as a reference when more information is required.
- The [Windows Server 2003 Technical Reference](http://www.microsoft.com/windows/reskits/default.asp) and the [Microsoft Windows Server 2003 Deployment Kit](http://www.microsoft.com/windows/reskits/default.asp) at www.microsoft.com/windows/reskits/default.asp.
- The “Deploying a Wireless LAN” chapter of the *Deploying Network Services* guide in the [Microsoft Windows Server 2003 Deployment Kit](http://www.microsoft.com/windowsserver2003/techinfo/reskit/deploykit.mspx) at www.microsoft.com/windowsserver2003/techinfo/reskit/deploykit.mspx.
This deployment kit chapter contains deployment guidance for using IAS in a number of scenarios that fall outside the scope of this secure wireless networking guidance, but that affect design decisions.
- For extensive coverage of 802.1X WLAN, WLAN security issues, and related standards, see [The Unofficial 802.11 Security Web Page](http://www.drizzle.com/~aboba/IEEE/) at www.drizzle.com/~aboba/IEEE/.
- For information about WLAN solutions and industry information, visit the [Wi-Fi Alliance](http://www.wi-fi-ally.com) Web site at www.wi-fi-ally.com.
- For information about WLAN, including background information, market research, white papers, and training programs, visit the [Wireless LAN Association \(WLANA\) Learning Center](http://www.wlana.org/learning_center.html) at www.wlana.org/learning_center.html.
- For information about EAP-TLS, EAPOL, EAP-RADIUS, RADIUS, and other Internet standards used with 802.1X, see [The Internet Engineering Task Force \(IETF\)](http://www.ietf.org/) Web site, at: www.ietf.org/.
- For information about relevant WLAN standards that include: 802.11, 802.11a, 802.11b, 802.1X, 802.11i, and others, see the [IEEE Wireless Standards Zone](http://standards.ieee.org/wireless/) Web site at <http://standards.ieee.org/wireless/>.

- The [802.11 Wireless Technical Reference](#) is available at www.microsoft.com/resources/documentation/windowsServ/2003/all/techref/en-us/W2K3TR_wir_intro.asp.
- For more information about 802.1X WLAN technologies, see the white paper, "[Windows XP Wireless Deployment Technology and Component Overview](#)" at www.microsoft.com/technet/itsolutions/desktopdeployment/mobility/wireless.mspx.