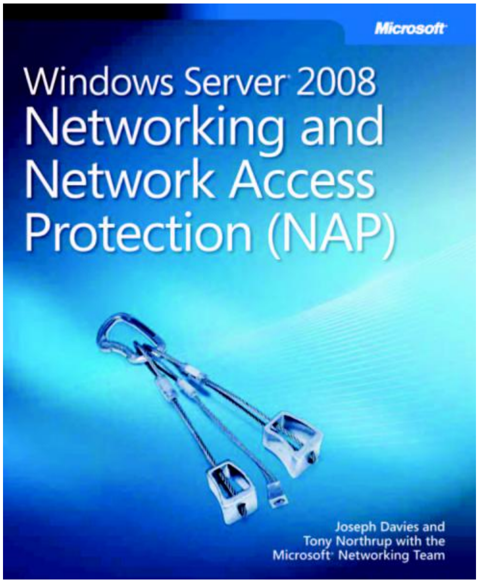

Windows Server® 2008 Networking and Network Access Protection (NAP)

Joseph Davies and Tony Northrup with the Microsoft Networking Team

To learn more about this book, visit Microsoft Learning at http://www.microsoft.com/MSPress/books/11160.aspx

9780735624221

Microsoft Press

Table of Contents

Part III Network Access Infrastructure

Part IV Network Access Protection Infrastructure

Chapter 9

Authentication Infrastructure

To deploy authenticated or protected network access, you must first deploy elements of a Microsoft Windows–based authentication infrastructure consisting of Active Directory, Group Policy, Remote Authentication Dial-In User Service (RADIUS), and a public key infrastructure (PKI). The set of elements you need to deploy depends on the type of network access and the design choices you make with regard to security, central configuration, and other issues. This chapter provides information about how to design and deploy these elements of an authentication infrastructure that can be used for wireless, wired, remote access, and site-to-site connections. Once deployed, elements of this infrastructure can also be used for Network Access Protection (NAP).

Concepts

The following sections provide technical background on the following technologies that are used in the Windows-based authentication infrastructure:

- Active Directory Domain Services
- Group Policy
- PKI
- RADIUS

Active Directory Domain Services

Active Directory Domain Services in the Windows Server 2008 operating system stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory uses a structured data store as the basis for a logical, hierarchical organization of directory information. Active Directory Domain Services can be installed on servers running Windows Server 2008.

This data store, or directory, contains Active Directory objects. These objects typically include shared resources such as servers, volumes, printers, and the network user and computer accounts.

Security is integrated with Active Directory through logon authentication and through access control to objects in the directory. With a single network logon, administrators can manage and organize directory data throughout their network, and authorized users can access resources anywhere on the network. Policy-based administration eases the management of even the most complex network.

Active Directory also includes the following:

- A set of rules (or schema) that defines the classes of objects and attributes contained in the directory, the constraints and limits on instances of these objects, and the format of their names.
- A global catalog that contains information about every object in the directory. This catalog allows users and administrators to find directory information regardless of which domain in the directory actually contains the data.
- A query and index mechanism, which enables objects and their properties to be published and found by network users or applications.
- A replication service that distributes directory data across a network. All domain controllers in a domain participate in replication and contain a complete copy of all directory information for their domain. Any change to directory data is replicated to all domain controllers in the domain.

User Accounts

Active Directory user accounts and computer accounts represent a physical entity such as a person, computer, or device. User accounts can also be used as dedicated service accounts for some applications.

User accounts and computer accounts (and groups) are also referred to as security principals. *Security principals* are directory objects that are automatically assigned security identifiers (SIDs), which can be used to access domain resources. A user or computer account is used to do the following:

- **Authenticate the identity of a user or computer.** A user account in Active Directory enables a user to log on to computers and domains with an identity that can be authenticated by the domain. Each user who logs on to the network should have his or her own unique user account and password. To maximize security, you should avoid multiple users sharing one account.
- **Authorize or deny access to domain resources.** When the user is authenticated, the user is authorized or denied access to domain resources based on the explicit permissions assigned to that user on the resource.
- **Administer other security principals.** Active Directory creates a foreign security principal object in the local domain to represent each security principal from a trusted external domain.
- **Audit actions performed using the user or computer account.** Auditing can help you monitor account security.

You can manage user or computer accounts by using the Active Directory Users And Computers snap-in.

Each computer that is running the Windows Vista, Windows XP, Windows Server 2008, or Windows Server 2003 operating system and that participates in a domain has an associated computer account. Similar to user accounts, computer accounts provide a means for authenticating and auditing computer access to the network and to domain resources.

User and computer accounts can be added, disabled, reset, and deleted using the Active Directory Users And Computers snap-in. A computer account can also be created when you join a computer to a domain.

Dial-In Properties of an Account

User and computer accounts in Active Directory contain a set of dial-in properties that can be used when allowing or denying a connection attempt. In an Active Directory–based domain, you can set the dial-in properties on the Dial-In tab of the user and computer account properties dialog box in the Active Directory Users And Computers snap-in. Figure 9-1 shows the Dial-In tab for a user account in a Windows Server 2008 functional level domain.

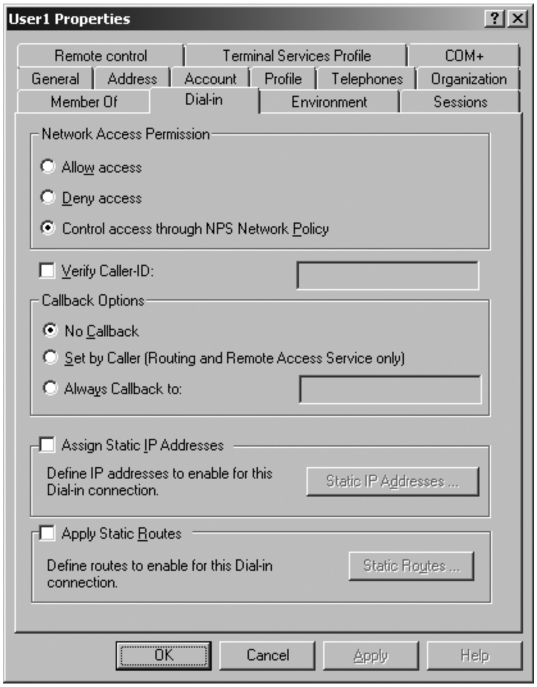

Network Access Permission

○ Allow access

○ Deny access

● Control access through NPS Network Policy

□ Verify Caller-ID:

Callback Options

● No Callback

○ Set by Caller (Routing and Remote Access Service only)

○ Always Callback to:

□ Assign Static IP Addresses

Define IP addresses to enable for this Dial-in connection.

□ Apply Static Routes

Define routes to enable for this Dial-in connection.

Figure 9-1 The Dial-In tab of a user account properties dialog box in a Windows Server 2008 functional level domain

On the Dial-In tab, you can view and configure the following properties:

- **Network Access Permission** You can use this property to set network access permission to be explicitly allowed, denied, or determined through Network Policy Server (NPS) network policies. NPS network policies are also used to authorize the connection attempt. If access is explicitly allowed, NPS network policy conditions and settings and account properties can still deny the connection attempt. The Control Access Through NPS Network Policy option is available on user and computer accounts in a Windows Server 2008 functional level domain. New accounts that are created for a Windows Server 2008 functional level domain are set to Control Access Through NPS Network Policy.

- **Verify Caller ID** If this property is enabled, the access server verifies the caller's phone number. If the caller's phone number does not match the configured phone number, the connection attempt is denied. This setting is designed for dial-in connections.
- **Callback Options** If this property is enabled, the access server calls the caller back during the connection process. Either the caller or the network administrator sets the phone number that is used by the server. This setting is designed for dial-in connections.
- **Assign Static IP Addresses** You can use this property to assign a specific IP address to a user when a connection is made. This setting is designed for dial-in connections.
- **Apply Static Routes** You can use this property to define a series of static IP routes that are added to the routing table of the server running the Routing and Remote Access service when a connection is made. This setting is designed for demand-dial routing.

Groups

A *group* is a collection of user and computer accounts and other groups that can be managed as a single unit. Users and computers that belong to a particular group are referred to as group members. Using groups can simplify administration by assigning a common set of permissions and rights to many accounts at once rather than assigning permissions and rights to each account individually.

Groups can be either directory-based or local to a particular computer. Active Directory provides a set of default groups upon installation and also allows you to create groups.

Groups in Active Directory allow you to do the following:

- Simplify administration by assigning permissions on a shared resource to a group rather than to individual users. This assigns the same access on the resource to all members of that group.
- Delegate administration by assigning user rights once to a group through Group Policy and then adding to the group members who require the same rights as the group.

Groups have a scope and type. Group *scope* determines the extent to which the group is applied within a domain or forest. Active Directory defines universal, global, and domain local scopes for groups. Group *type* determines whether a group can be used to assign permissions to a shared resource (for security groups); it also determines whether a group can be used for e-mail distribution lists only (for distribution groups).

Nesting allows you to add a group as a member of another group. You nest groups to consolidate member accounts and reduce replication traffic. Nesting options depend on the functional level of your domain. There are usually multiple domain functional levels, allowing for a phased upgrade of an environment, enabling additional domain-native functionality at each progressive level.

When you have decided how to nest groups based on your domain functional level, organize your user and computer accounts into the appropriate logical groups for the organization. For a Windows Server 2008 functional level domain, you can use universal and nested global groups. For example, create a universal group named WirelessUsers that contains global groups of wireless user and computer accounts for wireless intranet access. When you configure your NPS network policy for wireless access, you must specify only the WirelessUsers group name.

> **More Info** For more information about the types of groups, group scope, and domain functional levels, see the *Windows Server 2008 Active Directory Resource Kit* (Microsoft Press, 2008), which is available both as a stand-alone title and in the *Windows Server 2008 Resource Kit* (Microsoft Press, 2008); Windows Server 2008 Help and Support; or the resources at *http://www.microsoft.com/ad*.

Public Key Infrastructure

A *public key infrastructure* (PKI) is a system of digital certificates and certification authorities (CAs) that verifies and authenticates the validity of each entity—such as a user, computer, or Windows service—that is participating in secure communications through the use of public key cryptography.

Certification Authorities

When a certificate is presented to an entity as a means of identifying the certificate holder (the subject of the certificate), it is useful only if the entity being presented the certificate trusts the issuing CA. When you trust an issuing CA, it means that you have confidence that the CA has the proper policies in place when evaluating certificate requests and will deny certificates to any entity that does not meet those policies. In addition, you trust that the issuing CA will revoke certificates that should no longer be considered valid and will publish an up-to-date certificate revocation list (CRL). For more information about CRLs, see "Certificate Revocation" later in this chapter.

For Windows users, computers, and services, trust in a CA is established when you have a copy of the self-signed certificate of the root CA of the issuing CA locally installed and there is a valid certification path to the issuing CA. For a certification path to be valid, there cannot be any certificates in the certification path that have been revoked or whose validity periods have expired. The certification path includes every certificate issued to each CA in the certification hierarchy from a subordinate issuing CA to the root CA. For example, for a root CA, the certification path consists of a single certificate: its own self-signed certificate. For a subordinate CA, just below the root CA in the hierarchy, its certification path consists of two certificates: its own certificate and the root CA certificate.

If your organization is using Active Directory, trust in your organization's certification authorities will typically be established automatically based on decisions and settings made during the PKI deployment. For example, when joining a domain, a computer will automatically receive the organization's root CA through Group Policy settings.

Certification Hierarchies

A certification hierarchy provides scalability, ease of administration, and consistency with a growing number of commercial and other CA products. In its simplest form, a certification hierarchy consists of a single CA. However, in general, a hierarchy will contain multiple CAs with clearly defined parent-child relationships. In this model, the subordinate certification authorities are certified by their parent CA–issued certificates, which bind a CA's public key to its identity. The CA at the top of a hierarchy is referred to as the *root authority*, or *root CA*. The child CAs of the root CAs are called *subordinate CAs*.

In Windows Server 2008 and Windows Vista, if you trust a root CA (when you have its certificate in your Trusted Root Certification Authorities certificate store), you trust every subordinate CA in the hierarchy unless a subordinate CA has had its certificate revoked by the issuing CA or has an expired certificate. Thus, any root CA is an important point of trust in an organization and should be secured and maintained accordingly.

Verification of certificates thus requires trust in only a small number of root CAs. At the same time, it provides flexibility in the number of certificate-issuing subordinate CAs. There are several practical reasons for supporting multiple subordinate CAs, including the following:

- **Usage** Certificates can be issued for a number of purposes, such as securing e-mail and network authentication. The issuing policy for these uses can be distinct, and separation provides a basis for administering these policies.
- **Organizational divisions** There might be different policies for issuing certificates, depending upon an entity's role in the organization. You can create subordinate CAs for the purpose of separating and administering these policies.
- **Geographic divisions** Organizations might have entities at multiple physical sites. Network connectivity between these sites might dictate a requirement for multiple subordinate CAs to meet usability requirements.
- **Load balancing** If your PKI will support the issuing of a large number of certificates, having only one CA issue and manage all these certificates can result in considerable network load for that single CA. Using multiple subordinate certification authorities to issue the same kind of certificates divides the network load among certification authorities.
- **Backup and fault tolerance** Multiple certification authorities increase the possibility that your network will always have operational certification authorities available to service users.

Such a certificate hierarchy also provides administrative benefits, including the following:

- Flexible configuration of the CA security environment to tailor the balance between security and usability.

  For example, you might choose to employ special-purpose cryptographic hardware on a root CA, operate it in a physically secure area, or operate it offline. These security measures might be unacceptable for subordinate CAs because of cost or usability considerations.

- The ability to deactivate a specific portion of the CA hierarchy without affecting the established trust relationships.

  For example, you can easily shut down and revoke an issuing CA certificate that is associated with a specific geographic site without affecting other parts of the organization.

By using the Certificates snap-in, you can view the certification path for a certificate on the Certification Path tab of the properties dialog box of a certificate.

For a small business environment, a certificate hierarchy consisting of a single root CA that is also the issuing CA is adequate. For a medium-sized organization, a single root CA with a single level of issuing CAs is adequate. For an enterprise network, you can deploy a three-tiered CA hierarchy, consisting of the following:

- A root CA that is offline (not available on the network)
- A layer of intermediate CAs that are offline
- A layer of issuing CAs that are online

This CA hierarchy provides flexibility and insulates the root CA from attempts by malicious users to compromise its private key. The offline root and intermediate CAs are not required to be Windows Server 2008–based or Windows Server 2003–based CAs. Issuing CAs can be subordinates of a third-party intermediate CA. Figure 9-2 shows a three-level enterprise network certificate hierarchy.

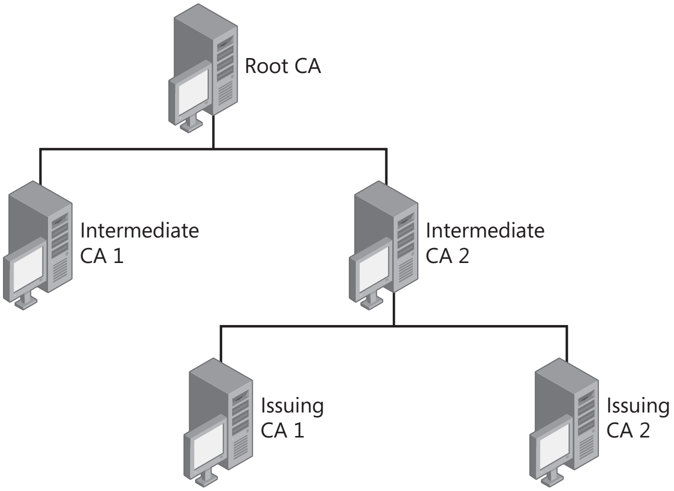

Figure 9-2 Three-level certificate hierarchy for enterprise networks

Certificate Revocation

Revocation of a certificate invalidates that certificate as a trusted security credential prior to the natural expiration of its validity period. There are a number of reasons why a certificate, as a security credential, could become untrustworthy prior to its expiration, including the following:

- Compromise or suspected compromise of the certificate subject's private key
- Compromise or suspected compromise of a CA's private key
- Discovery that a certificate was obtained fraudulently
- Change in the status of the certificate subject as a trusted entity
- Change in the name of the certificate subject

A PKI depends on distributed verification of credentials in which there is no need for direct communication with the central trusted entity that vouches for the credentials. This creates a need to distribute certificate revocation information to individuals, computers, and applications attempting to verify the validity of certificates. The need for revocation information and its timeliness will vary according to the application and its implementation of certificate revocation checking. To effectively support certificate revocation, the validating entity must determine whether the certificate is valid or has been revoked.

Certificate revocation lists (CRLs) are digitally signed lists of unexpired certificates that have been revoked. Clients retrieve this list and can then cache it (based on the configured lifetime of the CRL) and use it to verify certificates presented for use. Because CRLs can become large, depending on the size of the CA, delta CRLs can also be published. *Delta CRLs* contain only the certificates revoked since the last base CRL was published, which allows clients to retrieve the smaller delta CRL and quickly build a complete list of revoked certificates. The use of delta CRLs also allows more frequent publishing because the size of the delta CRL usually does not require as much overhead as a full CRL.

Windows Server 2008 supports industry-standard methods of certificate revocation. These methods include publication of CRLs and delta CRLs in several locations for clients to access in Active Directory and on Web servers and network file shares. Certificate revocation also can be checked by using the Online Certificate Status Protocol (OCSP), which uses the Hypertext Transfer Protocol (HTTP) to obtain a definitive digitally signed response indicating a certificate's revocation status.

Certificate Validation

The certificates that are offered during the negotiation for secure communication must be validated before secure communication can begin. For example, for network access authentication using Extensible Authentication Protocol-Transport Layer Security (EAP-TLS), the authentication server (the RADIUS server) must validate the certificate offered by the IEEE

802.1X or Point-to-Point Protocol (PPP) client. For authentication using either EAP-TLS or Protected EAP (PEAP), the 802.1X or PPP client can be configured to validate the certificate offered by the authentication server.

Windows Certificate Support

Windows has built-in support for certificates as follows:

- Every computer running Windows Vista, Windows Server 2008, Windows XP, or Windows Server 2003 has the ability, subject to Windows security and permissions, to store computer and user certificates and manage them by using the Certificates snap-in.
- Windows Server 2008 includes Active Directory Certificate Services and Windows Server 2003 includes Certificate Services, both of which allow a Windows server to act as a CA.

Certificate Services provides customizable services for issuing and managing certificates used in software security systems employing public key technologies. You can use Certificate Services in Windows Server 2008 and Windows Server 2003 to create a CA that will receive certificate requests, verify both the information in the request and the identity of the requester, issue certificates, revoke certificates, and publish CRLs.

You can also use Certificate Services to do the following:

- Enroll users for certificates from the CA by using a Web page (known as Web enrollment), through the Certificates snap-in, or transparently through autoenrollment.
- Use certificate templates to help simplify the choices that a certificate requester must make when requesting a certificate, depending upon the policy used by the CA.
- Take advantage of Active Directory for publishing trusted root certificates to domain member computers, publishing issued certificates, and publishing CRLs.
- Implement the ability to log on to a Windows domain by using a smart card.

If your organization is using Certificate Services, the CA is one of two types:

- **Enterprise CA** An enterprise CA depends on Active Directory being present. An enterprise CA offers different types of certificates to a requester based on the certificates it is configured to issue in addition to the security permissions of the requester. An enterprise CA uses information available in Active Directory to help verify the requester's identity. An enterprise CA can publish its CRL to Active Directory, a Web site, or a shared directory. You can use the Certificate Request Wizard within the Certificates snap-in, CA Web pages (Web enrollment), and autoenrollment to request certificates from an enterprise CA.
- **Standalone CA** For a user, a Standalone CA is less automated than an enterprise CA because it does not require or depend on the use of Active Directory. Standalone certification authorities that do not use Active Directory generally must request that the

certificate requester provide more complete identifying information. A Standalone CA makes its CRL available from a shared folder or from Active Directory if it is available. By default, users can request certificates from a Standalone CA only through Web enrollment.

> **More Info** For more information about PKI support in Windows, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/pki*.

Group Policy

The Group Policy management solution in Windows allows administrators to set configurations for both server and client computers. Local policy settings can be applied to all computers, and for those that are part of a domain, an administrator can use Group Policy to set policies that apply across a given site, domain, or organizational unit (OU) in Active Directory or that apply to a security group. Support for Group Policy is available on computers running Windows Vista, Windows Server 2008, Windows XP, and Windows Server 2003.

Through an Active Directory infrastructure and Group Policy, administrators can take advantage of policy-based management to do the following:

- Enable one-to-many management of users and computers throughout the enterprise.
- Automate enforcement of IT policies.
- Simplify administrative tasks such as system updates and application installations.
- Consistently implement security settings across the enterprise.
- Efficiently implement standard computing environments for groups of users.

Group Policy can be used to specify user-related policies and security, networking, and other policies applied at the computer level for management of domain controllers, member servers, and desktop user computers.

The GPMC snap-in provides a unified graphical user interface for deploying and managing Group Policy settings and enables script-based management of Group Policy operations. You can also use the Group Policy Management Editor snap-in.

On Windows Server 2008, you must install the Group Policy Management feature to use the Group Policy management tools such as the GPMC snap-in and Group Policy Management Editor snap-in.

Group Policy Overview

Administrators can manage computers centrally through Active Directory and Group Policy. Using Group Policy to deliver managed computing environments allows administrators to work more efficiently because of the centralized, one-to-many management it enables. Measurements of total cost of ownership (TCO) associated with administering distributed personal computer networks reveal lost productivity for users as one of the major costs for corporations. Lost productivity is frequently attributed to user errors—such as modifying system configuration files and thus rendering a computer unusable—or to complexity, such as the availability of nonessential applications and features on the desktop. Because Group Policy defines the settings and allowed actions for users and computers, it can create desktops that are tailored to users' job responsibilities and level of experience with computers.

Setting Group Policy By creating Group Policy settings, administrators use Group Policy to specify configurations for groups of users and computers. These settings are specified through the GPMC snap-in or the Group Policy Management Editor snap-in and are contained in a Group Policy Object (GPO), which is in turn linked to Active Directory containers—such as sites, domains, and OUs—and security groups.

In this way, Group Policy settings are applied to the users and computers in those Active Directory containers or security groups. Administrators can configure the users' work environment once and rely on the user's computer to enforce the policies as set.

Group Policy Capabilities Through Group Policy, administrators set the policies that determine how applications and operating systems are configured to keep users and systems functional and secure. Group Policies can be used for the following:

- **Registry-based policy** The most common and the easiest way to provide a policy for an application or operating system component is to implement a registry-based policy. By using the GPMC snap-in or the Group Policy Management Editor snap-in, administrators can create registry-based policies for applications, the operating system, and its components. For example, an administrator can enable a policy setting that removes the Run command from the Start menu for all affected users.
- **Security settings** Group Policy provides to administrators options for setting security options for computers and users within the scope of a GPO. Local computer, domain, and network security settings can be specified. For added protection, you can apply software restriction policies that prevent users from running files based on the path, URL zone, hash, or publisher criteria. You can make exceptions to this default security level by creating rules for specific software.

Using Group Policy

Administrators use Group Policy and Active Directory together to institute policies across domains, sites, and OUs according to the following rules:

- GPOs are stored on a per-domain basis.
- Multiple GPOs can be associated with a single site, domain, or OU.
- Multiple sites, domains, or OUs can use a single GPO.

- Any site, domain, or OU can be associated with any GPO, even across domains (although doing so slows performance).
- The effect of a GPO can be filtered to target particular groups of users or computers based on their membership in a security group.

Computer and User Configuration Administrators can configure specific desktop environments and enforce policy settings on groups of computers and users on the network as follows:

- **Computer configuration** Computer-related policies specify operating system behavior, desktop behavior, application settings, security settings, assigned applications options, and computer startup and shutdown scripts. Computer-related policy settings are applied during the computer startup process and during a periodic refresh of Group Policy.
- **User configuration** User-related policies specify operating system behavior, desktop settings, application settings, security settings, assigned and published application options, user logon and logoff scripts, and folder redirection options. User-related policy settings are applied when users log on to the computer and during the periodic refresh of Group Policy.

Applying Group Policy Group Policy is applied in an inherited and cumulative fashion and affects all computers and users in an Active Directory container. Group Policy is applied when the computer starts up and when the user logs on. When a user turns on the computer, the system applies computer-based Group Policy settings. When a user logs on interactively, the system loads the user's profile and then applies user-based Group Policy settings. By default, policy settings are reapplied every 90 minutes. (You can set this period between 0 and 45 days.) You can also locally reapply policy settings on demand by running the **gpupdate** command at a Windows command prompt.

When applying policy, the system queries the directory service for a list of GPOs to process. Active Directory resources that are enforced with Group Policy settings will require read access to the GPOs. If a computer or user is not allowed access to a GPO, the system does not apply the specified policy settings. If access is permitted, the system applies the policy settings specified by the GPO.

The scope of Group Policy can extend from a single computer—the local GPO that all computers include—to Active Directory sites, domains, and OUs. For example, a GPO might be linked to an Active Directory site to specify policy settings for proxy settings and network-related settings that are specific to that site. A GPO becomes useful only after it is linked to a container—the settings in the GPO are then applied according to the scope of the container.

GPOs are processed in the order of local, site, domain, and then OU. As a result, a computer or user receives the policy settings of the last Active Directory container processed—that is, a policy applied later overwrites policy applied earlier.

More Info For more information about Group Policy in Windows, see the *Microsoft Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/gp*.

RADIUS

When deploying a network access authentication infrastructure, it is possible to have each network access server store the account information and credentials for authentication and the network access policies for connection authorization. When a connection attempt is made, the access server can authenticate the connection attempt against the locally stored accounts and credentials, evaluate whether the connection attempt is authorized through the local account properties and network access policies, and locally store information about the connection attempt for later analysis. However, this method does not scale, especially in an enterprise environment with a large number of access servers. A scalable and more manageable solution is to offload the authentication and authorization evaluation and the storage of each connection attempt onto a central server that can utilize the existing accounts database.

RADIUS is a widely deployed protocol that allows authentication, authorization, and accounting for network access to be centralized at RADIUS servers. Originally developed for dial-up remote access, RADIUS is now supported by wireless access points (APs), authenticating Ethernet switches, virtual private network (VPN) servers, Digital Subscriber Line (DSL) access servers, and other types of network access servers.

More Info RADIUS is described in Request for Comments (RFC) 2865, "Remote Authentication Dial-In User Service (RADIUS)," and RFC 2866, "RADIUS Accounting." The listed RFCs can be viewed at *http://www.ietf.org/rfc.html*.

Components of a RADIUS Infrastructure

A RADIUS authentication, authorization, and accounting infrastructure consists of the following components:

- Access clients
- Access servers (RADIUS clients)
- RADIUS servers
- User account databases
- RADIUS proxies

Figure 9-3 shows the components of a RADIUS infrastructure.

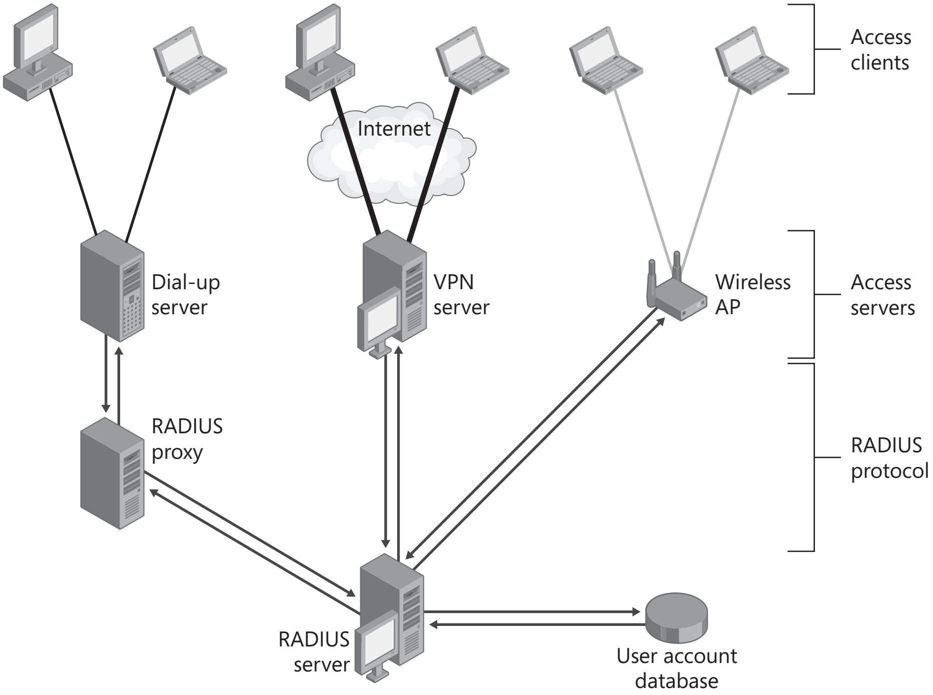

Figure 9-3 The components of a RADIUS infrastructure

These components are described in detail in the following sections.

Access Clients An access client requires access to a network or another part of the network. Examples of access clients are dial-up or VPN remote access clients, wireless clients, or LAN clients connected to an authenticating switch. Access clients are not RADIUS clients.

Access Servers (RADIUS Clients) An access server provides access to a network. An access server using a RADIUS infrastructure is also a RADIUS client, which uses the RADIUS protocol to send connection requests and accounting messages to a RADIUS server. Examples of access servers include:

- Wireless APs that provide physical layer access to an organization's network by using wireless-based transmission and reception technologies.
- Switches that provide physical layer access to an organization's network by using traditional LAN technologies such as Ethernet.
- Network access servers (NASs) that provide remote access connectivity to an organization's network or the Internet. An example is a computer running Windows Server 2008 and Routing and Remote Access and providing either traditional dial-up access or VPN-based remote access to an organization's intranet.

- Network Access Protection (NAP) enforcement points that collect a NAP client's system health status and send it to a Windows Server 2008–based RADIUS server for evaluation. Examples include NAP-enabled Dynamic Host Configuration Protocol (DHCP) servers and Health Registration Authorities (HRAs). For more information about NAP enforcement points, see Chapter 14, "Network Access Protection Overview."

RADIUS Servers A RADIUS server receives and processes connection requests or accounting messages sent by RADIUS clients or RADIUS proxies. During a connection request, the RADIUS server processes the list of RADIUS attributes in the connection request. Based on a set of rules and the information in the user account database, the RADIUS server authenticates and authorizes the connection and sends back either an accept or reject message. The accept message can contain connection restrictions that are enforced by the access server for the duration of the connection.

Note The NPS component of Windows Server 2008 is an industry standard–compliant RADIUS server.

User Account Databases A user account database is a list of user accounts and their properties that can be checked by a RADIUS server to verify authentication credentials and to obtain user account properties containing authorization and connection setting information.

The two user account databases that NPS can use are the local Security Accounts Manager (SAM) and Active Directory. For Active Directory, NPS can provide authentication and authorization for user or computer accounts in the domain in which the NPS server is a member, two-way trusted domains, and trusted forests with domain controllers running Windows Server 2008 or Windows Server 2003.

If the user accounts for authentication reside in a different type of database, you can use a RADIUS proxy to forward the authentication request to another RADIUS server that has access to the user account database.

RADIUS Proxies A RADIUS proxy routes RADIUS connection requests and accounting messages between RADIUS clients and RADIUS servers. The RADIUS proxy uses information within the RADIUS message to route the RADIUS message to the appropriate RADIUS client or server.

A RADIUS proxy can be used as a forwarding point for RADIUS messages when the authentication, authorization, and accounting must occur at multiple RADIUS servers within an organization or in different organizations.

With the RADIUS proxy, the definitions of *RADIUS client* and *RADIUS server* become blurred. A RADIUS client to a RADIUS proxy can be an access server (that originates connection requests or accounting messages) or another RADIUS proxy (in a chained proxy configuration). There can be multiple RADIUS proxies between the originating RADIUS client and the final RADIUS server using chained RADIUS proxies. In a similar way, a RADIUS server to a RADIUS proxy can be the final RADIUS server (at which the RADIUS message is evaluated for authentication and authorization) or another RADIUS proxy. Therefore, when referring to RADIUS clients and servers from a RADIUS proxy perspective, a RADIUS client is the RADIUS entity that receives RADIUS request messages, and a RADIUS server is the RADIUS entity that forwards RADIUS request messages.

Note The NPS component of Windows Server 2008 is an industry standard–compliant RADIUS proxy.

How It Works: RADIUS Messages and the RADIUS Authentication, Authorization, and Accounting Process

RADIUS messages are sent as User Datagram Protocol (UDP) messages. RADIUS authentication messages are sent to destination UDP port 1812, and RADIUS accounting messages are sent to UDP port 1813. Legacy access servers might use UDP port 1645 for RADIUS authentication messages and UDP port 1646 for RADIUS accounting messages. Only one RADIUS message is included in the UDP payload of a RADIUS packet.

A RADIUS message consists of a RADIUS header and RADIUS attributes. Each RADIUS attribute contains a specific item of information about the connection. For example, there are RADIUS attributes for the user name, the user password, the type of service requested by the user, the type of access server, and the IP address of the access server.

RADIUS attributes are used to convey information between RADIUS clients, RADIUS proxies, and RADIUS servers. For example, the list of attributes in the RADIUS Access-Request message includes information about the user credentials and the parameters of the connection attempt. In contrast, the list of attributes in the Access-Accept message includes information about the type of connection that can be made, connection constraints, and any vendor-specific attributes (VSAs).

More Info RADIUS attributes are described in RFCs 2548, 2865, 2866, 2867, 2868, 2869, 3162, and 3579. RFCs and Internet drafts for VSAs define additional RADIUS attributes. The listed RFCs can be viewed at *http://www.ietf.org/rfc.html*.

RFCs 2865 and 2866 define the following RADIUS message types:

- **Access-Request** Sent by a RADIUS client to request authentication and authorization for a network access connection attempt.
- **Access-Challenge** Sent by a RADIUS server in response to an Access-Request message. This message is a challenge to the RADIUS client that requires a

response. The Access-Challenge message is typically used for challenge-response based authentication protocols to verify the identity of the access client.

- **Access-Accept** Sent by a RADIUS server in response to an Access-Request message. This message informs the RADIUS client that the connection attempt is authenticated and authorized.
- **Access-Reject** Sent by a RADIUS server in response to an Access-Request message. This message informs the RADIUS client that the connection attempt is rejected. A RADIUS server sends this message if the credentials are not authentic or if the connection attempt is not authorized.
- **Accounting-Request** Sent by a RADIUS client to specify accounting information for a connection that was accepted.
- **Accounting-Response** Sent by the RADIUS server in response to the Accounting-Request message. This message acknowledges the successful receipt and processing of the Accounting-Request message.

For PPP authentication protocols such as Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), and Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), the results of the authentication negotiation between the access server and the access client are forwarded to the RADIUS server for verification in the Access-Request message.

For EAP–based authentication, the negotiation occurs between the RADIUS server and the access client. The RADIUS server uses Access-Challenge messages to send EAP messages to the access client. The access server forwards EAP messages sent by the access client to the RADIUS server as Access-Request messages. Within the Access-Challenge and Access-Request messages, EAP messages are encapsulated as the *RADIUS EAP-Message* attribute.

Authentication, authorization, and accounting of network access connections typically use RADIUS messages in the following way. (See Figure 9-3.)

1. Access servers—such as dial-up network access servers, VPN servers, and wireless APs—receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the RADIUS server.
3. The RADIUS server evaluates the Access-Request message.
4. If required (for example, when the authentication protocol is EAP), the RADIUS server sends an Access-Challenge message to the access server. The response to the challenge is sent as a new Access-Request to the RADIUS server. This can occur multiple times during the EAP negotiation.

5. The RADIUS server verifies the user credentials and the authorization of the connection attempt.
6. If the connection attempt is both authenticated and authorized, the RADIUS server sends an Access-Accept message to the access server. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends an Access-Reject message to the access server.
7. Upon receipt of the Access-Accept message, the access server completes the connection process with the access client and sends an Accounting-Request message to the RADIUS server.
8. After the Accounting-Request message is processed, the RADIUS server sends an Accounting-Response message.

Planning and Design Considerations

The following sections describe key planning and design considerations for the following technologies in a Windows-based network access authentication infrastructure:

- Active Directory
- PKI
- Group Policy
- RADIUS

Active Directory

It is beyond the scope of this book to describe in detail the planning and design considerations for deploying Active Directory in an organization of arbitrary size. For detailed information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or resources at *http://www.microsoft.com/ad*.

The following sections describe the planning and design considerations for Active Directory that will help you create a manageable Windows-based authentication infrastructure for network access.

Accounts and Groups

Depending on the type of connection, network access authentication can use the credentials and properties of user or computer accounts. For each type, you must ensure that the Network Access Permission on the Dial-In tab is set to either Allow Access or Control Access Through NPS Network Policy (recommended). By default, new computer and user accounts have the Network Access Permission set to Control Access Through NPS Network Policy.

Accounts contain the account name and an encrypted form of the account password that can be used for validation of the client's credentials. Additional account properties determine whether the account is enabled or disabled, locked out, or permitted to log on only during specific hours. If an account is disabled, locked out, or not permitted to log on during the time of the connection, the connection attempt is rejected.

When using groups to manage access, you can use your existing groups and create network policies in NPS that either allow access (with or without restrictions) or reject access based on the group name. For example, you can configure an NPS network policy that specifies the Employees group, which has no network access restrictions for VPN connections. You can also configure another network policy that specifies that the accounts in the Contractors group can create VPN connections only during business hours.

NPS can use Active Directory user principal names (UPNs) and universal groups. In a large domain with thousands of users, create a universal group for all of the users for whom you want to allow access, and then create a network policy that grants access for this universal group. To minimize the processing of group membership for a user account, do not put all of your user accounts directly into the universal group, especially if you have a large number of user accounts. Instead, create separate global groups that are members of the universal group, and add user accounts to those global groups.

Domain and Forest Trust Relationships

The NPS server is an Active Directory domain member and can verify authentication credentials for accounts in the domain of which it is a member and in all other domains that trust the NPS server's domain. Therefore, ensure that all of the domains in your Active Directory infrastructure trust the domain of the NPS server (subject to security restrictions and policies for your organization); otherwise, you must configure the NPS server as a RADIUS proxy to forward the connection request messages to another NPS server that can authenticate the user or computer account that is attempting to connect.

For the NPS server to be able to access the dial-in properties for user and computer accounts, you must add the computer account of the NPS server to the RAS and IAS Servers group for each domain: the domain of the NPS server and all the domains that trust the NPS server's domain.

PKI

It is beyond the scope of this book to describe in detail the planning and design considerations for deploying a PKI in an organization of arbitrary size. For detailed information, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/pki*.

A PKI is needed for the following purposes in a Windows-based network access infrastructure:

- Autoenrollment of computer certificates on domain member computers for computer-level certificate-based network access
- Autoenrollment of user certificates on domain member computers for user-level certificate-based network access
- Automatic provisioning of computer health certificates on domain member computers for Internet Protocol security (IPsec) enforcement when deploying NAP.

Subsequent chapters in this book describe additional PKI requirements for different types of network access and for NAP.

The following planning and design considerations for your PKI are specific to a Windows-based authentication infrastructure for network access:

- When using certificates for computer-level network access authentication, configure Group Policy for autoenrollment of computer certificates.

  Examples are the use of EAP-TLS or Protected EAP-TLS (PEAP-TLS) for computer-level wireless authentication.
- When using certificates for user-level network access authentication, configure a certificate template for user certificates, and configure Group Policy for autoenrollment of user certificates.

  Examples are the use of EAP-TLS or PEAP-TLS for user-level wireless authentication.
- When using PEAP-MS-CHAP v2 for network access authentication, configure Group Policy for autoenrollment of computer certificates to install computer certificates on the NPS servers. You can use computer certificates when NPS is not installed on an Active Directory domain controller. Alternatively, you can use the RAS and IAS Server certificate template and configure autoenrollment for members of the RAS and IAS Servers security group.

  Examples are the use of PEAP-MS-CHAP v2 for computer-level or user-level wireless authentication.
- When using IPsec enforcement in NAP, you might need to configure a certificate template for health certificates.
- When using certificates for computer-level or user-level network access authentication, ensure that the CRLs are published in a primary location and in at least one secondary location and that these locations are accessible by all computers, especially the RADIUS servers. The RADIUS servers will first attempt to validate the certificate by using OSCP. If the OSCP validation is not successful, the RADIUS server will attempt to perform a CRL validation of the user or computer certificate. By default, the NPS RADIUS servers

will reject all certificate-based connection attempts if they cannot verify the certificate's revocation status.

Direct from the Source: Modifying CLR Checking Behavior

Performing CRL checking is enabled by default for security reasons. It is possible to modify the behavior of NPS for certificate revocation checking. There are special cases in which you might want or need to make this change; three examples are as follows:

- If your PKI environment has a poor or slow CRL distribution infrastructure
- If you are using third-party certificates that do not or are not able to provide CRL distribution points with the most up-to-date CRLs
- If you rely on an external distribution point and do not have redundant external connections

Any of these conditions could lead to problems with the certificate revocation checking, thus causing delays or intermittent authentication failure. If you must modify NPS for your deployment, you will be making changes to values in the following registry key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13

The two values you will be most concerned with are:

- **IgnoreNoRevocationCheck** This is set to 0 by default. When set to 1, NPS allows the clients to connect even when it does not perform or cannot complete a revocation check.
- **NoRevocationCheck** This is set to 0 by default. When set to 1, NPS does not attempt a revocation check.

If you set either or both of these registry keys to 1, simply revoking someone's certificate won't limit their network access.

Chris Irwin, Premier Field Engineer

Premier Field Engineering Group

Group Policy

It is beyond the scope of this book to describe in detail the planning and design consideration for deploying Group Policy in an organization of arbitrary size. For detailed information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/gp*.

Group Policy is used for the following purposes in a Windows-based network access authentication infrastructure:

- To deploy settings to install a root certificate on domain member computers in order to validate the computer certificates of the NPS servers
- To deploy settings to autoenroll computer certificates on domain member computers for computer-level certificate-based network access authentication
- To deploy settings to autoenroll user certificates on domain member computers for user-level certificate-based network access authentication

Additionally, Group Policy allows you to deploy configuration settings for the following:

- IEEE 802.11 wireless network profiles
- Wired (Ethernet with 802.1X authentication) network profiles
- Windows Firewall with Advanced Security connection security rules to protect traffic
- NAP client configuration

When planning your Group Policy infrastructure, adhere to the recommendations and best practices for Group Policy configuration within your Active Directory infrastructure, as described in the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or in the resources at *http://www.microsoft.com/gp*. There are no specific planning and design considerations for Group Policy objects that are specific to a Windows-based authentication infrastructure for network access and for NAP. However, you must ensure that the correct Group Policy Objects are being applied to those containers or security groups that contain user or computer accounts for authenticated access or for configuration of wireless or wired network profiles, Windows Firewall with Advanced Security connection security rules, or NAP client settings.

RADIUS

NPS can be used as a RADIUS server, a RADIUS proxy, or both. The following sections describe the planning, design, and security considerations when deploying NPS as a RADIUS server or proxy.

RADIUS Server Planning and Design Considerations

When planning to deploy an NPS-based RADIUS infrastructure for network access authentication or for NAP, consider the following:

- **Domain membership for NPS servers** You must determine the domain in which to make the NPS server a member. For multiple domain environments, an NPS server can authenticate credentials for user accounts in the domain of which it is a member and all domains that trust its domain. To read the dial-in properties for user and computer accounts, however, you must add the computer account of the NPS server to the RAS and IAS Servers groups for each domain.

- **UDP ports for RADIUS traffic** If needed, you can configure the NPS server to receive RADIUS messages that are sent to UDP ports other than the default ports of 1812 and 1645 (for RADIUS authentication) and ports 1813 and 1646 (for RADIUS accounting).
- **RADIUS clients to configure on the NPS server** A RADIUS client can be an access server—a network access server (for example, a dial-up or VPN server, a wireless AP, or an Ethernet switch) or a NAP enforcement point—or a RADIUS proxy. NPS supports all access servers and RADIUS proxies that comply with RFC 2865. Configure each access server or RADIUS proxy that sends RADIUS request messages to the NPS server as a RADIUS client on the NPS server.

  You can specify IP addresses or DNS names for RADIUS clients. In most cases, it is better to specify IPv4 or IPv6 addresses for RADIUS clients. When you use IP addresses, NPS is not required to resolve host names at startup and will start much more quickly. This is beneficial especially if your network contains a large number of RADIUS clients. Use DNS names to specify RADIUS clients when you require something other than administrative flexibility (for example, the ability to map multiple RADIUS client addresses to a single DNS name).

  NPS in Windows Server 2008 allows you to specify a RADIUS client by using an address range. The address range for IPv4-based RADIUS clients is expressed in the network prefix length notation *w.x.y.z/p*, where *w.x.y.z* is the dotted decimal notation of the address prefix, and *p* is the prefix length (the number of high order bits that define the network prefix). This is also known as Classless Inter-Domain Routing (CIDR) notation. An example is 192.168.21.0/24. To convert from subnet mask notation to network prefix length notation, *p* is the number of high order bits in the subnet mask that are set to 1. The address range for IPv6-based RADIUS clients is also expressed in network prefix length notation. An example is 2001:db8:27a1:1c5d::/64.
- **Wireless APs, switches, and third-party remote access servers** To determine whether a third-party access server is interoperable with NPS as a RADIUS server, refer to the third-party access server documentation for its RFC 2865 compliance and its use of RADIUS attributes and vendor-specific attributes.
- **Connection request policy configuration** Connection request policies determine whether the NPS server is used as a RADIUS server, a RADIUS proxy, or both, depending on the information in the incoming RADIUS request messages. The Use Windows Authentication For All Users default connection request policy is configured for NPS when it is used as a RADIUS server. Additional connection request policies can be used to specify more specific conditions, manipulate attributes, and specify advanced attributes. Connection request policies are processed in order, so place the more specific policies at the top of the list. You use the Network Policy Server snap-in to manage new connection request policies.

- **Realm replacement to convert user name formats** The *realm* name is the part of the account name that identifies the location of the user account, such as the name of an Active Directory domain. To correctly replace or convert realm names within the user name of a connection request, configure realm name rules for the User-Name RADIUS attribute on the appropriate connection request policy.
- **Network policy configuration** Network policies are used to grant or deny network access and to set specific conditions for allowed network access, such as dial-in constraints, allowed authentication protocols and encryption strength, and additional RADIUS attributes. Use the Network Policy Server snap-in to manage network policies.
- **Network policies and authorization by user or group** In small organizations, you can manage authorization by setting the network access permission on each user account. For a large organization, set the network access permission on each user account to be controlled through the settings of an NPS network policy. Then, configure network policies to grant access by using group membership.
- **Additional RADIUS attributes and vendor-specific attributes** If you plan to return additional RADIUS attributes or vendor-specific attributes (VSAs) with the responses to RADIUS requests, you must add the RADIUS attributes or VSAs to the appropriate network policy.
- **Event logging** Event logging for authentication events, enabled by default, can assist with troubleshooting connection attempts.
- **Access logging** Access logging stores the authentication and accounting request messages received from access servers and collects this information in a central location. You can store the information in local log files or a Microsoft SQL Server database.
- **Interim accounting** Some access servers send interim accounting messages periodically during a connection, in contrast to the accounting message that is sent when the connection attempt is made. To use interim accounting, first verify that your access server supports sending interim accounting messages. Next, add the Acct-Interim-Interval RADIUS attribute as a standard RADIUS attribute from the Settings tab of the appropriate network policy. Configure the Acct-Interim-Interval attribute with the interval (in minutes) to send periodic interim accounting messages.

RADIUS Server Security Considerations

When using NPS as a RADIUS server, consider the following to ensure a protected RADIUS infrastructure:

- **RADIUS shared secrets** RADIUS shared secrets are used to verify that RADIUS messages, with the exception of the Access-Request message, are sent by a RADIUS-enabled device that is configured with the same shared secret. Shared secrets also verify that the RADIUS message has not been modified in transit (message integrity). The shared secret is also used to encrypt some sensitive RADIUS attributes, such as User-Password and Tunnel-Password. Configure strong shared secrets and change them frequently to prevent dictionary attacks. Strong shared secrets are a long (more than 22 characters) sequence of random letters, numbers, and punctuation. You can use the Network Policy Server snap-in to generate strong RADIUS shared secrets.

- **Message Authenticator attribute** To ensure that an incoming RADIUS Access-Request message—for connection requests that use the PAP, CHAP, MS-CHAP, and MS-CHAP v2 authentication protocols—was sent from a RADIUS client configured with the correct shared secret, you can use the RADIUS Message Authenticator attribute (also known as a *digital signature* or the *signature attribute*). You must enable the use of the Message Authenticator attribute on both the NPS server (as part of the configuration of the RADIUS client in the Network Policy Server snap-in) and the RADIUS client (the access server or RADIUS proxy). Ensure that the RADIUS client supports the Message Authenticator attribute before enabling it. The Message Authenticator attribute is always used with EAP-based authentication methods.

  For information about enabling the RADIUS Message Authenticator attribute for your access server, see your access server documentation.

- **Firewall configuration for RADIUS traffic** If your NPS server is on a perimeter network, configure your Internet firewall (between your perimeter network and the Internet) to allow RADIUS traffic to pass between your NPS server and RADIUS clients on the Internet. You might need to configure an additional firewall that is placed between your perimeter network and your intranet to allow traffic to flow between the NPS server on the perimeter network and domain controllers on the intranet.
- **Network access authentication protocols** NPS includes support for several different authentication protocols. The order of included authentication protocols, from the most secure to the least secure, is: PEAP-TLS, EAP-TLS, PEAP-MS-CHAP v2, MS-CHAP v2, CHAP, and PAP. Microsoft recommends using only the strongest authentication protocols that are required for your configuration. For password-based authentication protocols, strong password policies must be enforced to protect from dictionary attacks. The use of PAP is not recommended unless it is required.

Direct from the Source: EAP-MD5 Removed

With the release of Windows Vista, the Microsoft EAP-MD5 implementation has been removed. The decision to remove the Microsoft EAP-MD5 implementation was made in the interest of improving security in Windows Vista. The removal of the Microsoft implementation of EAP-MD5 directly affects remote access services, VPN services, and wired 802.1X deployments. By default, these components can no longer use the Microsoft EAP-MD5 implementation for authentication. The server implementation of EAP-MD5 will continue to ship with Windows Server 2008, but it will be disabled by default. Microsoft will continue to terminate EAP-MD5 connections for legacy network devices but will not initiate them from Microsoft's client operating systems.

Tim Quinn, Support Escalation Engineer

Enterprise Platform Support

- **Remote access account lockout** To provide protection for online dictionary attacks launched against access servers by using known user names, you can enable remote access account lockout. Remote access account lockout disables remote access for user accounts after a configured number of failed connection attempts has been reached. For more information, see Chapter 12, "Remote Access VPN Connections."

  Remote access account lockout can also be used to prevent a malicious user from intentionally locking out a domain account by attempting multiple dial-up or VPN connections with the wrong password. You can set the number of failed attempts for remote access account lockout to a number that is lower than the logon retries for domain account lockout. By doing this, remote access account lockout occurs before domain account lockout, which prevents the domain account from being intentionally locked out.
- **Certificates to install on NPS servers for network access authentication** When you use the included EAP-TLS, PEAP-TLS, or PEAP-MS-CHAP v2 authentication protocols, by default you must install a computer certificate on the NPS server containing the Server Authentication purpose in the Enhanced Key Usage (EKU) extensions. Other authentication protocols provided by independent software or hardware vendors might also require certificates on NPS servers.
- **Using Windows Firewall with Advanced Security connection security rules to protect NPS servers** You can configure Windows Firewall with Advanced Security connection security rules to protect RADIUS traffic sent between RADIUS servers and access servers and between RADIUS servers and RADIUS proxies with IPsec. These rules can be configured as part of Group Policy settings and applied to Active Directory containers or filtered for security groups, or they can be created and applied to individual servers.

RADIUS Proxy Planning and Design Considerations

When planning to deploy a RADIUS infrastructure for network access authentication or for NAP, consider the following:

- **When to use NPS as a RADIUS proxy** The following uses of NPS as a RADIUS proxy are described in this chapter:
  - ❑ When you want to provide authentication and authorization for user accounts that are not members of either the domain in which the NPS server is a member or another domain that has a two-way trust with the domain in which the NPS server is a member. This includes accounts in untrusted domains, one-way trusted domains, and other forests. Instead of configuring your access servers to send their connection requests to an NPS RADIUS server, you can configure them to send their connection requests to an NPS RADIUS proxy. The NPS RADIUS proxy uses the realm name portion of the user name to forward the request to an NPS server in the correct domain or forest. Connection attempts for user accounts in one domain or forest can be authenticated for network access servers that are members of another domain or forest.

  - When you want to process a large number of connection requests. In this case, instead of configuring your RADIUS clients to attempt to balance their connection and accounting requests across multiple RADIUS servers, you can configure them to send their connection and accounting requests to an NPS RADIUS proxy. The NPS RADIUS proxy dynamically balances the load of connection and accounting requests across multiple RADIUS servers and increases the processing of large numbers of RADIUS clients and authentications per second.

  For more information about these configurations, see "Using RADIUS Proxies for Cross-Forest Authentication" and "Using RADIUS Proxies to Scale Authentications" later in this chapter.

- **Connection request policy configuration** The Use Windows Authentication For All Users default connection request policy uses NPS as a RADIUS server. To create a connection request policy to use NPS as a RADIUS proxy, you must first create a remote RADIUS server group whose members are the set of RADIUS servers to which a RADIUS message is forwarded. Next, create a connection request policy that forwards authentication requests to a remote RADIUS server group. Finally, either delete the Use Windows Authentication For All Users connection request policy or move the new connection request policy higher in the list so that it is evaluated first.

- **Realm replacement and attribute manipulation** To convert realm names and configure RADIUS message forwarding based on the realm name, you must use realm rules for the User-Name attribute on the appropriate connection request policy. If you are using the MS-CHAP v2 authentication protocol, you cannot manipulate the User Name attribute if the connection request policy is used to forward the RADIUS message. The only exception occurs when a backslash character (\) is used, and the manipulation affects only the information to the left of it. A backslash character is typically used to indicate a domain name (the information to the left of the backslash) and a user account name within the domain (the information to the right of the backslash). In this case, only attribute manipulation rules that modify or replace the domain name are allowed.

- **The use of additional RADIUS attributes and vendor-specific attributes** If you plan to include additional RADIUS attributes and vendor-specific attributes (VSAs) to RADIUS requests that are being forwarded, you must add the RADIUS attributes and VSAs to the appropriate connection request policy.

- **Remote RADIUS server group configuration** A remote RADIUS server group contains the set of RADIUS servers to which RADIUS messages matching a connection request policy are forwarded.

- **Copying logging information at the NPS proxy** The NPS proxy can record all RADIUS accounting information that it receives in the local log file. This creates a central location for all authentication and accounting information for all of the access servers of the NPS proxy.

- **Authentication and accounting ports** When you configure a server in a remote RADIUS server group, you can configure custom UDP ports to which RADIUS authentication and accounting messages are sent. The default UDP port for authentication requests is 1812. The default UDP port for accounting requests is 1813.
- **Load balancing and failure detection** When you configure multiple servers in a remote RADIUS server group, you can configure settings that determine how the NPS proxy balances the load of authentication and accounting requests over the RADIUS servers in the group. By default, the RADIUS traffic is balanced equally across the members of the group. You can use additional settings to configure NPS to detect and recover from the failure of a remote RADIUS server group member.

Direct from the Source: RADIUS Proxies and Trusts

It is best to avoid creating arbitrary trusts for cross-domain network authentication. If your goal is to allow domain users the ability to log on to networks in different domains, use RADIUS proxies rather than a transitive trust. With a RADIUS proxy, you are passing only the essential data between the two NPS servers necessary for granting user or computer access. Additionally, this requires at most only two UDP ports to be available between the two domains. With a trust, far more traffic, such as resource access validation, is being passed, and many more ports are required to be opened.

Clay Seymour, Support Escalation Engineer

Enterprise Platform Support

RADIUS Proxy Security Considerations

When using NPS as a RADIUS proxy, consider the following to ensure a protected RADIUS infrastructure:

- **Shared secrets** Configure strong shared secrets to prevent dictionary attacks, and change them frequently. Strong shared secrets are a long (more than 22 characters) sequence of random letters, numbers, and punctuation.
- **Firewall configuration** If your NPS proxy is on a perimeter network, configure your Internet firewall (between your perimeter network and the Internet) to allow RADIUS messages to pass between your NPS proxy and RADIUS clients on the Internet. You might need to configure an additional firewall that is placed between your perimeter network and your intranet to allow RADIUS traffic to flow between the NPS proxy on the perimeter network and an NPS server on the intranet.
- **Message Authenticator attribute** You can use the RADIUS Message Authenticator attribute (also known as a *digital signature* or the *signature attribute*) to ensure that RADIUS Access-Request messages for connection requests were sent from a RADIUS

client configured with the correct shared secret. The Message Authenticator attribute is always used with EAP, and you don't have to enable it on the NPS server or access server. For the PAP, CHAP, MS-CHAP, and MS-CHAP v2 authentication protocols, you must enable the use of the Message Authenticator attribute on both the NPS server (as part of the configuration of the RADIUS client) and the RADIUS client (the access server or RADIUS proxy). Ensure that the RADIUS client supports the Message Authenticator attribute before enabling it.

- **Using Windows Firewall with Advanced Security connection security rules to protect NPS proxies** You can configure the Windows Firewall with Advanced Security connection security rules to use IPsec to protect RADIUS traffic sent between NPS proxies and access servers and between the NPS proxies and RADIUS servers.
- **Password Authentication Protocol (PAP)** The use of PAP is strongly discouraged, especially when using RADIUS proxies.

High Availability for RADIUS Authentication

To provide high availability for RADIUS-based authentication and accounting, you should always use at least two NPS servers. One NPS server is used as the primary RADIUS server, and the other is used as a backup. Access servers or other RADIUS proxies are configured for both NPS servers (a primary and a secondary) and automatically switch to the secondary NPS RADIUS server when the primary NPS RADIUS server becomes unavailable. When using multiple RADIUS servers, failover is based on a RADIUS client switching to another RADIUS server and performing a new authentication transaction. Failover within a transaction is not supported.

High Scalability for RADIUS Authentication

Consider the following for scaling RADIUS authentication to an organization containing a large number of accounts or connection attempt activity:

- **Use universal groups and group-based network policies** If you are using network policies to restrict access for all but certain groups, create a universal group for all of the users or computers for whom you want to allow access, and then create a network policy that grants access for this universal group. Do not put all of your user and computer accounts directly into the universal group, especially if you have a large number of them on your network. Instead, create separate groups that are members of the universal group, and add the user and computer accounts to those groups.
- **Use user principal names** Use user principal names (UPNs), such as user@contoso.com, to refer to users whenever possible. A user can have the same user principal name regardless of domain membership. This practice provides scalability that might be required in organizations with a large number of domains.

- **Install NPS on domain controllers** If possible, install NPS on domain controllers for best authentication and authorization performance. When NPS is running on a domain controller, the traffic and processing delays incurred when an NPS RADIUS server contacts a domain controller over the network to verify account credentials and obtain account properties are eliminated.

  If the NPS server is on a computer other than a domain controller, and it is receiving a large number of authentication requests per second, you can improve performance by increasing the number of concurrent authentications between the NPS server and the domain controller. To do this, edit the following registry key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters. Add a new value (REG_DWORD value type) named MaxConcurrentApi, and although the range can be between 0 and 10, assign it a setting from 2 through 5.

  This value specifies the maximum number of simultaneous logon calls that can be transmitted to the domain controller over the secure channel at any given time, and the default is 2 for a member server computer. Increasing the setting will allow additional logon calls to be processed simultaneously to improve performance on the NPS server. Avoid setting the MaxConcurrentApi value to a setting higher than 5 because the additional load might cause depletion of resources on the domain controller.

Deployment Steps

This section contains the steps or resources for the steps to deploy the following components of a Windows-based network access authentication infrastructure:

- Active Directory
- PKI
- Group Policy
- RADIUS

Deploying Active Directory

It is beyond the scope of this book to instruct you on the specific steps to deploy Active Directory for an organization of arbitrary size. For additional information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/ad*.

The elements of configuring Active Directory to best support a Windows-based authentication infrastructure for network access are as follows:

- Ensure that all users who are making user-level authenticated connections have a corresponding user account that is enabled.

- Ensure that all computers that are making computer-level authenticated connections have a corresponding computer account that is enabled.
- Set the network access permission on user and computer accounts to the appropriate setting: either Allow Access or Control Access Through NPS Network Policy (recommended). The network access permission setting is on the Dial-In tab on the properties dialog box of a user or computer account in the Active Directory Users And Computers snap-in.
- Organize your network access user and computer accounts into the appropriate groups. Use a Windows 2000, Windows Server 2003, or Windows Server 2008 functional-level domain and universal groups and global groups to organize your accounts for a specific type of access into a single group. For example, for wireless access, create a universal group named WirelessUsers that contains global groups of wireless user and computer accounts for intranet access.

Deploying PKI

It is beyond the scope of this book to provide the specific steps to deploy a PKI for an organization of arbitrary size. For additional information, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/pki*.

The elements of configuring a certificate services-based PKI to best support a Windows-based authentication infrastructure for network access are as follows:

- When using certificates for user-level network access authentication, configure a certificate template for user certificates. If you are running a Windows enterprise CA, you can make a copy of the standard user template. Standalone CAs do not support certificate templates.
- When using IPsec enforcement in NAP, you might need to configure a certificate template for health certificates.

> **Note** A certificate template for a computer certificate is already configured by default with Windows Certificate Services.

After your PKI has been deployed, there are a set of procedures for deploying certificates that are common to wireless, wired, remote access VPN, and site-to-site VPN connections. These procedures are as follows:

- Configuring autoenrollment of computer certificates to computers in an Active Directory domain
- Using the Certificates snap-in to request a computer certificate

- Using the Certificates snap-in to import a computer certificate
- Executing a CAPICOM script that requests a computer or user certificate
- Configuring autoenrollment of user certificates in an Active Directory domain
- Using the Certificates snap-in to request a user certificate
- Using the Certificates snap-in to import a user certificate
- Installing third-party certificate chains by using Group Policy
- Requesting a certificate via the Web

Configuring the Autoenrollment of Computer Certificates to Computers in an Active Directory Domain If you are using a Windows Server 2008 enterprise CA as an issuing CA, each computer can automatically request a computer certificate from the issuing CA by using a Computer Configuration group policy setting. This method allows a single point of configuration for an entire domain.

To Configure an Active Directory Domain for Automatic Enrollment of Computer Certificates

1. Open the Group Policy Management snap-in.
2. In the console tree, expand Forest, expand Domains, and then click the name of the domain to which your CA belongs.
3. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
4. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, then Windows Settings, then Security Settings, and then Public Key Policies.
5. Right-click Automatic Certificate Request Settings, point to New, and then click Automatic Certificate Request.
6. The Automatic Certificate Request Setup Wizard appears. Click Next.
7. On the Certificate Template page, click Computer, and then click Next.
8. Click Finish.

To immediately obtain an updated Computer Configuration Group Policy to request a computer certificate for a computer running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP, restart the computer, or type **gpupdate /target:computer** at a command prompt.

Using the Certificates Snap-In to Request a Computer Certificate If you are using a Windows Server 2008 enterprise CA as an issuing CA, each computer can separately request a computer certificate from the issuing CA by using the Certificates snap-in.

To Request a Computer Certificate by Using the Certificates Snap-In

1. Log on to the computer using an account that has administrator privileges for that computer.
2. On the Start menu, click Run, type **mmc**, and then press Enter.
3. On the Console menu, click File, and then click Add/Remove Snap-In.
4. In the Add Or Remove Snap-Ins dialog box, under Available Snap-Ins, double-click Certificates. In the Certificates Snap-In dialog box, click Computer Account, and then click Next.
5. Do one of the following:
   - ❑ To manage certificates for the local computer, click Local Computer.
   - ❑ To manage certificates for a remote computer, click Another Computer and type the name of the computer, or click Browse to select the computer name. Then click OK.
6. Click Finish. Certificates (Local Computer) or Certificates (*computername*) appears on the list of selected snap-ins for the new console. Click OK.
7. In the console tree, expand the Certificates\Personal node.
8. Right-click the Personal node, point to All Tasks, and then click Request New Certificate.

The Certificate Request Wizard guides you through the steps of requesting a certificate. For a Windows-based client computer, the certificate imported into the Local Computer store must have the Client Authentication Enhanced Key Usage (EKU). For the certificate installed on the VPN server or the NPS server, the certificate imported into the Local Computer store must have the Server Authentication EKU.

Using the Certificates Snap-In to Import a Computer Certificate If you have a certificate file that contains the computer certificate, you can import the computer certificate by using the Certificates snap-in. This must be done when you purchase individual computer certificates for your VPN or RADIUS servers from a third-party CA for PEAP-MS-CHAP v2 authentication or for Secure Socket Tunneling Protocol (SSTP) connections.

To Import a Computer Certificate by Using the Certificates Snap-In

1. Open the Certificates (Local Computer)\Personal node.
2. Right-click the Personal node, point to All Tasks, and then click Import.

The Certificate Import Wizard guides you through the steps of importing a certificate from a certificate file. For a Windows-based client computer, the certificate imported into the Local Computer store must have the Client Authentication EKU. For the certificate installed on the VPN or NPS server, the certificate imported into the Local Computer store must have the Server Authentication EKU.

> **Note** It is also possible to import a certificate by double-clicking a certificate file that is stored in a folder or sent in an e-mail message. Although this works for certificates created with Windows-based CAs, this method might not work for third-party CAs. The recommended method of importing certificates is to use the Certificates snap-in.

Executing a CAPICOM Script That Requests a Computer or User Certificate In this method, each computer must execute a CAPICOM script that requests a computer or user certificate from the issuing CA. CAPICOM is a COM client that performs cryptographic functions (the CryptoAPI) by using Microsoft ActiveX and COM objects. CAPICOM can be used with Microsoft Visual Basic, Visual Basic Scripting Edition, and C++. For more information about CAPICOM, visit *http://msdn2.microsoft.com/en-us/library/ms995332.aspx*.

To perform an enterprise deployment of user and computer certificates, a CAPICOM program or script can be distributed through e-mail for execution, or users can be directed to a Web site containing a link to a CAPICOM program or script. Alternately, the CAPICOM program or script can be placed in the user's logon script file for automatic execution. The storage location of the user or computer certificate can be specified using the CAPICOM application programming interfaces (APIs).

Configuring Autoenrollment of User Certificates to Users in an Active Directory Domain This method allows a single point of configuration for the entire domain. All members of the domain automatically request the user certificate through a User Configuration group policy setting. If you use as an issuing CA an enterprise CA from Windows Server 2008, Windows Server 2003 Enterprise Edition, or Windows Server 2003 Datacenter Edition, you can install user certificates through autoenrollment.

To Configure User Certificate Enrollment for an Enterprise CA

1. On the Start menu, click Run, type **mmc**, and then click OK.
2. On the File menu, click Add/Remove Snap-In.
3. Under Available Snap-Ins, double-click Certificate Templates, and then click OK.
4. In the console tree, click Certificate Templates. All certificate templates appear in the details pane.
5. In the details pane, right-click the User template, and then click Duplicate Template. When prompted for the minimum version of the CA to support the certificate template, click Windows Server 2003, Enterprise Edition, and then click OK.
6. In the Template Display Name field, type the name of the new user certificate template (for example, **VPNAccess**).

   Make sure that the Publish Certificate In Active Directory check box is selected.
7. Click the Security tab.

8. In the Group Or User Names list, click Domain Users.
9. In the Permissions For Domain Users list, select the Read, Enroll, and Autoenroll permission check boxes, and then click OK.
10. Open the Certification Authority snap-in.
11. In the console tree, expand your CA's name, and then click Certificate Templates.
12. On the Action menu, point to New, and then click Certificate Template To Issue.
13. Click the name of the newly created user certificate template (for example, VPNAccess), and then click OK.
14. Open the Group Policy Management snap-in.
15. In the console tree, expand Forest, expand Domains, and then click the name of your domain to which your CA belongs.
16. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
17. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, then Windows Settings, then Security Settings, and then Public Key Policies.
18. In the details pane, double-click Certificate Services Client – Auto–Enrollment.
19. In Configuration Model, select Enabled from the drop-down list.
20. Select the Renew Expired Certificates, Update Pending Certificates, and Remove Revoked Certificates check box.
21. Select the Update Certificates That Use Certificate Templates check box, and then click OK.

Perform steps 15–21 for each domain container, as appropriate. Ensure that all appropriate domain containers are configured for autoenrollment of user certificates, either through the inheritance of group policy settings of a parent container or through explicit configuration.

To immediately update User Configuration group policy and request a user certificate for a computer that is running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP and is a member of the domain for which autoenrollment is configured, restart the computer, or at a command prompt, type **gpupdate /target:user**.

Using the Certificates Snap-In to Request a User Certificate If you are using a Windows Server 2008 enterprise CA as an issuing CA, each computer can separately request a user certificate from the issuing CA by using the Certificates snap-in.

To Request a User Certificate by Using the Certificates Snap-In

1. Log on to the computer using an account that has administrator privileges for that computer.

2. On the Start menu, click Run, type **mmc**, and then press Enter.
3. On the Console menu, click File, and then click Add/Remove Snap-In.
4. In the Add Or Remove Snap-Ins dialog box, under Available Snap-Ins, double-click Certificates. In the Certificates Snap-In dialog box, click My User Account, click Finish, and then click OK.
5. In the console tree, expand the Certificates\Personal node.
6. Right-click the Personal node, point to All Tasks, and then click Request New Certificate.

The Certificate Request Wizard guides you through the steps of requesting a user certificate. For a Windows-based client computer, the imported certificate must have the Client Authentication EKU.

Using the Certificates Snap-In to Import a User Certificate If you have a certificate file that contains the user certificate, you can import the user certificate by using the Certificates snap-in.

To Import a User Certificate by Using the Certificates Snap-In

1. Open the Certificates (Current User)\Personal node.
2. Right-click the Personal node, point to All Tasks, and then click Import.

The Certificate Import Wizard guides you through the steps of importing a certificate from a certificate file. For a Windows-based client computer, the certificate imported into the Local Computer store must have the Client Authentication EKU.

Installing Third-Party Certificate Chains by Using Group Policy When you are using a third-party CA for the computer certificates that are installed on access servers or RADIUS servers, you might need to install the chain of certificates (the root CA certificate to the issuing CA certificate) for the certificate installed on the access or RADIUS server. If the access client does not trust the certificate chain of the certificate submitted by the access or RADIUS server, certificate validation can fail.

A certificate chain consists of the root CA certificate and the certificate of each intermediate CA, including the issuing CA. The following procedures describe how to deploy a root CA certificate and an intermediate CA certificate to access clients by using Group Policy.

To Install a Root CA Certificate by Using Group Policy

1. In the console tree of the Certificates snap-in for the access or RADIUS server computer account, expand Certificates (Local Computer), expand Trusted Root Certification Authorities, and then click Certificates.
2. In the details pane, right-click the root CA certificate of the issuing CA of the computer certificate on the authentication server, point to All Tasks, and then click Export.
3. In the Certificate Export Wizard, on the Welcome to the Certificate Export Wizard page, click Next.

4. On the Export File Format page, click Cryptographic Message Syntax Standard–PKCS #7 Certificates (.p7b).
5. Click Next. On the File To Export page, type the file name for the exported certificate, or click Browse to specify a location and file name.
6. Click Next. On the Completing The Certificate Export Wizard page, click Finish.
7. Open the Group Policy Management snap-in.
8. In the console tree, expand Forest, expand Domains, and then click the name of your domain to which your CA belongs.
9. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
10. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, Windows Settings, Security Settings, and then Public Key Policies.
11. Right-click Trusted Root Certification Authorities, and then click Import.
12. In the Certificate Import Wizard, specify the file that was saved in step 5.
13. Repeat steps 8–12 for all appropriate domain containers and their Group Policy Objects.

The next time the access client computers update their Computer Configuration group policy, the root CA certificates of the issuing CAs of the authentication server computer certificates are installed in their local certificate store.

To Install an Intermediate CA Certificate by Using Group Policy

1. In the console tree of the Certificates snap-in for the access or RADIUS server computer account, expand Certificates (Local Computer), expand Intermediate Certification Authorities, and then click Certificates.
2. In the details pane, right-click the intermediate CA certificate of the issuing CA of the computer certificate on the authentication server, point to All Tasks, and then click Export.
3. In the Certificate Export Wizard, on the Welcome To The Certificate Export Wizard page, click Next.
4. On the Export File Format page, click Cryptographic Message Syntax Standard–PKCS #7 Certificates (.p7b).
5. Click Next. On the File To Export page, type the file name for the exported certificate, or click Browse to specify a location and file name.
6. Click Next. On the Completing The Certificate Export Wizard page, click Finish.
7. Open the Group Policy Management snap-in.
8. In the console tree, expand Forest, expand Domains, and then click the name of your domain to which your CA belongs.

9. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
10. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, Windows Settings, Security Settings, and then Public Key Policies.
11. Right-click Intermediate Certification Authorities, point to All Tasks, and then click Import.
12. In the Certificate Import Wizard, specify the file that was saved in step 5.

Repeat steps 8–12 for all appropriate domain containers and their Group Policy Objects.

If you cannot use Group Policy, you can manually install root and intermediate certificates on individual access client computers.

To Manually Install a Root or Intermediate CA Certificate on an Access Client

1. Export the root CA certificate of the access or RADIUS server's computer certificate to a . p7b file.
2. On the access client computer, in the console tree of the Certificates (Local Computer) snap-in, expand Certificates (Local Computer), expand Trusted Root Certification Authorities (for a root CA certificate) or Intermediate Certification Authorities (for an intermediate CA certificate), and then click Certificates.
3. Right-click Certificates, point to All Tasks, and then click Import.
4. The Welcome To The Certificate Import Wizard page of the Certificate Import Wizard appears. Click Next.
5. On the File To Import page, in the File Name box, type the file name of the certificate file saved in step 1, or click Browse and use the Browse dialog box to locate it.
6. Click Next. On the Certificate Store page, click Place All Certificates In The Following Store, and then specify the import location.
7. Click Next. On the Completing The Certificate Import Wizard page, click Finish.

Requesting a Certificate via the Web Requesting a certificate via the Web, also known as Web enrollment, is done with Microsoft Windows Internet Explorer. For the address, type **http://*servername*/certsrv**, where *servername* is the computer name of the Windows Server 2008 or Windows Server 2003 CA that is also running Internet Information Services (IIS). A Web-based wizard takes you through the steps of requesting a certificate. The location where the certificate is stored (whether it is the Current User store or the Local Computer store) is determined by whether the Use Local Machine Store check box was selected when an advanced certificate request was performed. This check box is cleared by default, and certificates are stored in the Current User store. You must have local administrator privileges to store a certificate in the Local Computer store.

You can use Web enrollment with either an enterprise or a Standalone CA.

Direct from the Source: Duplicating Default Certificate Templates

When using certificate templates, you should always make a duplicate of the default template, and if applicable, make your scenario-specific changes to the new template. For example, if you want to change the security groups that can autoenroll for a user certificate, make a duplicate of the user certificate. Then, obtain the properties of the new certificate template, click the Security tab, and add the specific groups that you want to have access to the template.

Clay Seymour, Support Escalation Engineer

Enterprise Platform Support

Group Policy

It is beyond the scope of this book to provide the specific steps to deploy Group Policy for an organization of arbitrary size. For additional information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or resources at *http://www.microsoft.com/gp*.

The elements of configuring Group Policy to best support a Windows-based authentication infrastructure for network access are as follows:

- When using certificates for computer-level network access authentication, configure Group Policy for autoenrollment of computer certificates. This requires deployment of a Windows enterprise CA. Autoenrollment cannot be configured when using a Standalone CA.
- When using certificates for user-level network access authentication, configure a certificate template for user certificates, and configure Group Policy for autoenrollment of user certificates.
- When using PEAP-MS-CHAP v2 for network access authentication, optionally configure Group Policy for autoenrollment of computer certificates to install computer certificates on the NPS servers.
- When you are using PEAP-MS-CHAP v2 for network access authentication and a third-party CA for the computer certificates installed on the NPS RADIUS servers, ensure that the root CA certificate for the NAP RADIUS server's computer certificate is installed on the access clients. If not, configure Group Policy to install the appropriate root CA certificate on domain member computers.

For information about how to configure Group Policy to deploy certificate settings, see "Deploying PKI" earlier in this chapter.

For information about how to configure Group Policy to deploy configuration settings for specific types of network access, see the following:

- Chapter 10, "IEEE 802.11 Wireless Networks"
- Chapter 11, "IEEE 802.1X-Authenticated Wired Networks"
- Chapter 16, "IPsec Enforcement"
- Chapter 17, "802.1X Enforcement"
- Chapter 18, "VPN Enforcement"
- Chapter 19, "DHCP Enforcement"

RADIUS Servers

Configuring a fault-tolerant RADIUS infrastructure requires at a minimum the configuration of at least two NPS RADIUS servers, a primary NPS RADIUS server, and a secondary RADIUS NPS server. You must do the following:

- Configure the primary NPS server.
- Copy the configuration of the primary NPS server to the secondary NPS server.

Because the configuration of the primary NPS server is being copied to the secondary NPS server, you should always make configuration changes to the primary NSP server.

Configuring the Primary NPS Server

To configure the primary NPS server on a computer, complete these steps as discussed in the following sections:

1. Obtain and install a computer certificate.
2. Install NPS and configure NPS server properties.
3. Configure NPS with RADIUS clients.
4. Use IPsec to protect RADIUS traffic.
5. Configure the appropriate policies.

Obtaining and Installing a Computer Certificate If you have configured computer certificate autoenrollment, force a refresh of computer configuration Group Policy by typing **gpupdate /target:computer** at a command prompt.

If you use a Windows Server 2008 or Windows Server 2003 enterprise CA and you are not using autoenrollment for computer certificates, you can request one, as described in the following procedure.

To Request a Computer Certificate

1. Click Start, click Run, type **mmc**, and then click OK.
2. On the File menu, click Add/Remove Snap-In.
3. Under Available Snap-Ins, double-click Certificates, click Computer Account, and then click Next.
4. Do one of the following:
    - To manage certificates for the local computer, click Local Computer, and then click Finish.
    - To manage certificates for a remote computer, click Another Computer and type the name of the computer, or click Browse to select the computer name. Click Finish.
5. Click OK.
6. In the console tree, expand Certificates (Local Computer or *Computername*), and then click Personal.
7. On the Action menu, point to All Tasks, and then click Request New Certificate to start the Certificate Enrollment Wizard.
8. On the Before You Begin page, click Next.
9. On the Request Certificates page, click Computer, and then click Enroll.
10. Click Finish.

If your PKI does not support autoenrollment of computer certificates, obtain the computer certificate as a saved file, and then use the following procedure to import the computer certificate on the primary NPS server.

> **Note** To perform the next procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

To Import the Computer Certificate on the Primary NPS Server

1. In the console tree of the Certificates snap-in, expand Certificates (Local Computer or *Computername*).
2. Right-click Personal, point to All Tasks, and then click Import.
3. On the Welcome To The Certificate Import Wizard page, click Next.
4. On the File To Import page, in the File Name box, type the file name of the certificate file provided by the commercial CA. Alternatively, you can click Browse and use the Browse dialog box to locate it.

5. Click Next. On the Certificate Store page, click Place All Certificates In The Following Store. By default, the Personal node should appear as the import location. Click Next, and then click Finish.

Configuring NPS Server Properties NPS is installed on computers running Windows Server 2008 with the Network Policy and Access Services role through the Initial Configuration Tasks or Server Manager tools. However, the primary NPS server computer must be able to access account properties in the appropriate domains. If NPS is being installed on a domain controller, no additional configuration is required for NPS to access account properties in the domain to which it belongs. If NPS is not installed on a domain controller, you must configure the primary NPS server computer to read the properties of user accounts in the domain, as described in the following procedure:

To Configure the Primary NPS Server Computer to Read the Properties of User Accounts in the Domain

1. In the console tree of the Network Policy Server snap-in, right-click NPS (Local), and then click Register Server In Active Directory.
2. In the Network Policy Server dialog box, click OK twice.

Alternatively, you can do one of the following:

- Use the **netsh nps add registeredserver** command.
- Use the Active Directory Users And Computers snap-in to add the computer account of the NPS server to the RAS and IAS Servers security group.

If the NPS server authenticates and authorizes network access attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the NPS server computer is a member. Next, configure the NPS server computer to read the properties of user accounts in other domains by using the **netsh nps add registeredserver** command or by using the Active Directory Users And Computers snap-in.

If there are accounts in other domains, and the domains do not have a two-way trust with the domain in which the NPS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains. If there are accounts in other untrusted Active Directory forests, you must configure a RADIUS proxy between the forests. For more information, see “Using RADIUS Proxies for Cross-Forest Authentication” later in this chapter.

If you want to store authentication and accounting information for connection analysis and security investigation purposes, enable logging for accounting and authentication events. Windows Server 2008 NPS can log information to a local file and to a SQL Server database.

To Enable and Configure Local File Logging for NPS

1. In the console tree of the Network Policy Server snap-in, click Accounting.
2. In the details pane, click Configure Local File Logging.

3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests in the NPS log files:
   - To capture accounting requests and responses, select the Accounting Requests check box.
   - To capture authentication requests, access-accept packets, and access-reject packets, select the Authentication Requests check box.
   - To capture periodic status updates, such as interim accounting packets, select the Periodic Accounting Status or Periodic Authentication Status check boxes.

   All these logging options are enabled by default.
4. On the Log File tab, type the log file directory as needed, and then select the log file format and new log time period. The default log file directory is *%SystemRoot%*\System32\LogFiles.

To Enable and Configure SQL Server Database Logging for NPS

1. In the console tree of the Network Policy Server snap-in, click Accounting.
2. In the details pane, click Configure SQL Server Logging.
3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests. All these logging options are enabled by default.
4. In Maximum Number of Concurrent Sessions, type the maximum number of simultaneous sessions that NPS can create with SQL Server.
5. To configure a SQL data source, click Configure.
6. In the Data Link Properties dialog box, configure the appropriate settings for the SQL Server database.

If needed, configure additional UDP ports for authentication and accounting messages that are sent by RADIUS clients (the access servers). By default, NPS uses UDP ports 1812 and 1645 for authentication messages and UDP ports 1813 and 1646 for accounting messages.

To Configure NPS for Different UDP Ports

1. In the console tree of the Network Policy Server snap-in, right-click NPS, and then click Properties.
2. Click the Ports tab, and then in the Authentication section, type the UDP port numbers for your RADIUS authentication traffic. In the Accounting section, type the UDP port numbers for your RADIUS accounting traffic.

   To use multiple port settings for authentication or accounting traffic, separate the port numbers with commas. You can also specify an IP address to which the RADIUS messages must be sent by typing in the following syntax: ***IPAddress*:*UDPPort***. For example, if you have multiple network adapters and you want to receive RADIUS authentication messages sent only to the IP address of 10.0.0.99 and UDP port 1812, in the Authentication box, type **10.0.0.99:1812**. However, if you specify IP addresses and copy the configuration of the primary NPS server to the secondary NPS server, you must modify the ports on the secondary NPS server to either remove the IP address of the primary NPS server or change the IP address to that of the secondary NPS server.

Configuring NPS with RADIUS Clients You must configure the primary NPS server with the access servers or RADIUS proxies as RADIUS clients.

To Add a RADIUS Client for NPS

1. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients And Servers, right-click RADIUS Clients, and then click New RADIUS Client.
2. In the New RADIUS Client dialog box, under Name And Address, in the Friendly Name text box, type a name for the RADIUS client (the access server or RADIUS proxy). In the Address (IP Or DNS) text box, type the IP address or DNS domain name of the RADIUS client. If you type a DNS domain name, click Verify to resolve the name to the correct IP address for the access server.
3. Under Shared Secret, in the Shared Secret and Confirm Shared Secret text boxes, type the shared secret for this combination of NPS server and RADIUS client or click Generate to have the NPS service generate a strong RADIUS shared secret.
4. Under Additional Options, specify whether this RADIUS client will always use the Message-Authenticator attribute in RADIUS messages and whether the RADIUS client is a NAP enforcement point that is running Windows Server 2008 (the RADIUS Client Is NAP-Capable check box), and then click OK.

If you have multiple wireless APs on a single subnet, you can simplify RADIUS client administration by specifying an IPv4 or IPv6 address range instead of specifying the address or DNS name of a single RADIUS client. All of the RADIUS clients in the range must be configured to use the same RADIUS server and shared secret. If you are not using this feature, use a different shared secret for each wireless AP.

Use as many RADIUS shared secrets as you can. Each shared secret should be a random sequence of uppercase and lowercase letters, numbers, and punctuation marks that is at least 22 characters long. To create a strong RADIUS shared secret, use the Generate option when configuring a shared secret with the Network Policy Server snap-in.

Using IPsec to Protect RADIUS Traffic To ensure maximum security for RADIUS messages, it is recommended that you use IPsec and Encapsulating Security Payload (ESP) to provide data confidentiality, data integrity, and data origin authentication for RADIUS traffic sent between the NPS servers and the RADIUS clients. Computers running Windows Server 2008 and Windows Server 2003 support IPsec. You configure the NPS RADIUS server for IPsec protection of RADIUS traffic through Windows Firewall with Advanced Security connection security rules. To secure RADIUS traffic sent from third-party access servers, the access servers must also support IPsec. For more information about connection security rules, see Chapter 4, "Windows Firewall with Advanced Security."

Configuring the Appropriate Policies To evaluate authorization and connection constraints for incoming connection requests, you must configure the appropriate policies consisting of connection request policies, network policies, and for NAP, health policies. The Network Policy Server snap-in has a set of wizards to automatically configure a set of policies for common network access and NAP scenarios. The following procedure describes how to run the Network Policy Server wizards.

To Run the Network Policy Server Wizards

1. In the console tree of the Network Policy Server snap-in, click NPS (Local).
2. In the details pane, in the drop-down list select one of the following:
   - ❑ Network Access Protection (NAP)
   - ❑ RADIUS Server For Dial-up Or VPN Connections
   - ❑ RADIUS Server For 802.1X Wireless or Wired Connections
3. If you selected Network Access Protection (NAP), click Configure NAP and use the pages of the Configure NAP Wizard to specify the set of policies for NAP enforcement.
4. If you selected RADIUS Server For Dial-up Or VPN Connections, click Configure VPN Or Dial-up and use the pages of the Configure VPN Or Dial-up Wizard to specify the set of policies for VPN or dial-up–based network access.
5. If you selected RADIUS Server For 802.1X Wireless or Wired Connections, click Configure 802.1X and use the pages of the Configure 802.1X Wizard to specify the set of policies for VPN or dial-up–based network access.

See the following chapters for information about configuring the appropriate policies with the Network Policy Server wizards:

- Chapter 10
- Chapter 11
- Chapter 12
- Chapter 13, "Site-to-Site Connections"
- Chapter 16
- Chapter 17
- Chapter 18
- Chapter 19

If the access servers require vendor-specific attributes (VSAs), you must add the VSAs to the appropriate network policy.

To Add a VSA to a Network Policy

1. In the console tree of the Network Policy Server snap-in, expand Policies, and then click Network Policies.
2. Right-click the NPS network policy to which the VSA will be added, and then click Properties.
3. Click the Settings tab, click Vendor Specific, and then click Add. A list of predefined attributes appears in the Add Vendor Specific Attribute dialog box.
4. Look at the list of available RADIUS attributes to determine whether your vendor-specific attribute is already present. If it is, double-click it and configure it as specified in your access server's documentation.
5. If the vendor-specific attribute is not in the list of available RADIUS attributes, double-click Vendor-Specific. The Attribute Information dialog box appears.
6. Click Add. The Vendor-Specific Attribute Information dialog box appears.
7. To specify the network access server vendor for your access server from the list, click Select From List, and then select the network access vendor for which you are configuring the VSA.
8. If the vendor is not listed, click Enter Vendor Code, and then type the vendor code in the space provided.

> **More Info** If you do not know the vendor code for your access server, see RFC 1007 for a list of SMI Network Management Private Enterprise Codes. RFC 1007 can be viewed at *http://www.ietf.org/rfc.html.*

9. Specify whether the attribute conforms to the RFC 2865 VSA specification. If you are not sure, see your access server documentation. If your attribute conforms, click Yes. It Conforms, and then click Configure Attribute. The Configure VSA (RFC-Compliant) dialog box appears.
10. In the Vendor-Assigned Attribute Number spin box, type the number that is assigned to the attribute (the numbers available are 0 through 255). In the Attribute Format drop-down list, specify the format for the attribute, and then in the Attribute Value text box, type the value that you are assigning to the attribute. Click OK twice.
11. If the attribute does not conform, click No. It Does Not Conform, and then click Configure Attribute. The Configure VSA (Non-RFC-Compliant) dialog box appears.

12. In the Hexadecimal Attribute Value text box, type the value for the attribute. Click OK twice.

Configuring the Secondary NPS Server

To configure the secondary NPS server on a computer, do the following:

1. Obtain and install a computer certificate.
2. Configure the secondary NPS server computer to read the properties of user accounts in the domain.
3. Copy the configuration of the primary NPS server to the secondary NPS server.

Copying the Configuration of the Primary NPS Server to the Secondary NPS server To copy the configuration of the primary NPS server to the secondary NPS server, do the following:

1. On the primary NPS server computer, type **netsh nps export *path\file* exportpsk=yes** at a command prompt, which stores the configuration settings, including RADIUS shared secrets, in a text file at ***path\file***. The path can be a relative, an absolute, or a network path.
2. Copy the file created in step 1 to the secondary NPS server.
3. On the secondary NPS server computer, type **netsh nps import *path\file*** at a command prompt, which imports all the settings configured on the primary NPS server into the secondary NPS server.

If you must change the NPS server configuration in any way, use the Network Policy Server snap-in to change the configuration of the NPS server that is designated as the primary configuration server, and then use this procedure to synchronize those changes on the secondary NPS server.

Using RADIUS Proxies for Cross-Forest Authentication

Because NPS uses Active Directory to validate credentials and obtain user and computer account properties, a RADIUS proxy must be placed between the access servers and the NPS server computers when the user and computer accounts for access client computers and users exist in the following authentication databases:

- Two different Active Directory forests that do not trust each other
- Two different domains that do not trust each other
- Two different domains that have a one-way trust

Direct from the Source: RADIUS Proxies and EAP-TLS

The use of a RADIUS proxy is required for EAP-TLS because part of the process requires a service principal name (SPN) lookup in Active Directory. However, SPN lookups do not work across trusts. When the NPS server receives the computer identity, it is in the form of an SPN (*host/ComputerName.DNSDomainName*). The NPS server passes the SPN to the local global catalog. If the global catalog is unable to match the SPN to a local domain account, it will fail the request with a No Valid Account Found error condition. SPN requests are not passed to the other domains.

Clay Seymour, Support Escalation Engineer

Enterprise Platform Support

Note You do not need to use a RADIUS proxy if you use PEAP-MS-CHAP v2 and user names like those used prior to Windows 2000 (microsoft\user1, for example).

When an access client sends user credentials, a user name is often included, which includes two elements:

- Identification of the user account name
- Identification of the user account location

For example, for the user name user1@contoso.com, user1 is the user account name, and contoso.com is the location of the user account. The identification of the location of the user account is known as a *realm*, which has different forms:

- **The realm name can be a prefix.** In contoso\user1, *contoso* is the name of a domain like those used prior to Windows 2000.
- **The realm name can be a suffix.** For user1@contoso.com, *contoso.com* is either a DNS domain name or the name of an Active Directory–based domain.

The user name is passed from the access client to the access server during the authentication phase of the connection attempt. This user name becomes the User-Name RADIUS attribute in the Access-Request message sent by the access server to its configured RADIUS server, which is a RADIUS proxy in this configuration. When the RADIUS proxy receives the Access-Request message, connection request policies on the RADIUS proxy determine the RADIUS server to which the Access-Request message is forwarded based on the realm name.

Figure 9-4 shows NPS RADIUS proxies forwarding RADIUS messages between access servers and multiple NPS RADIUS servers in two different Active Directory forests.

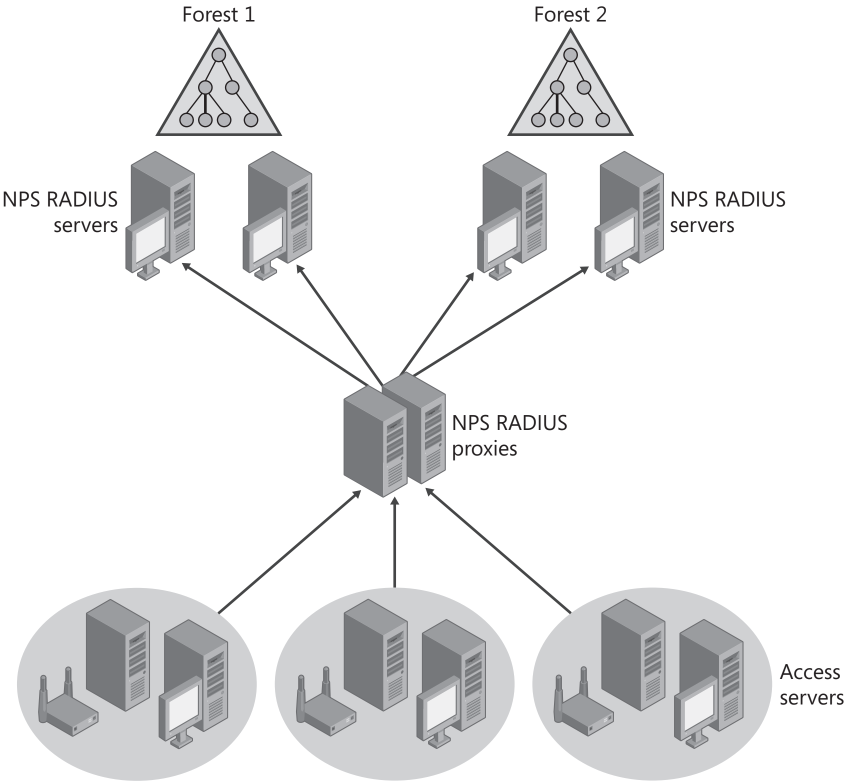

Figure 9-4 Using NPS RADIUS proxies for cross-forest authentication

The following configuration is for an organization that uses the following:

- **Active Directory domains** Active Directory domains contain the user accounts, passwords, and dial-in properties that each NPS RADIUS server requires to authenticate user credentials and evaluate authorization.
- **At least two NPS RADIUS servers in each forest** At least two NPS RADIUS servers (one primary and one secondary) can provide fault tolerance for RADIUS-based authentication, authorization, and accounting in each forest. If only one NPS RADIUS server is configured and it becomes unavailable, access clients for that forest cannot be authenticated. By using at least two NPS RADIUS servers and configuring the NPS RADIUS proxies for both the primary and secondary NPS RADIUS servers, the NPS RADIUS proxies can detect when the primary NPS RADIUS server is unavailable and then automatically fail over to the secondary NPS RADIUS server.
- **A network policy for network access** A network policy is configured on the NPS RADIUS servers to authorize network connections based on group membership.
- **At least two NPS RADIUS proxies** At least two NPS RADIUS proxies can provide fault tolerance for RADIUS requests that are sent from the access servers.

To deploy the configuration just described, do the following:

1. Configure the certificate infrastructure.
2. Configure the Active Directory forests for accounts and groups.
3. Configure the primary NPS RADIUS server on a computer in the first forest.
4. Configure the secondary NPS RADIUS server on another computer in the first forest.
5. Configure the primary NPS RADIUS server on a computer in the second forest.
6. Configure the secondary NPS RADIUS server on another computer in the second forest.
7. Configure the primary NPS RADIUS proxy.
8. Configure the secondary NPS RADIUS proxy.
9. Configure RADIUS authentication and accounting on the access servers.

Configuring the Certificate Infrastructure

Follow the instructions in the “Deploying PKI” subsection of “Deployment Steps” earlier in this chapter.

Configuring the Active Directory Forests for Accounts and Groups

Follow the instructions in the “Deploying Active Directory” subsection of “Deployment Steps” earlier in this chapter.

Configuring the Primary NPS Server on a Computer in the First Forest

To configure the primary NPS RADIUS server on a computer in the first forest, perform on a computer in the first forest the steps described in the following subsections of “Configuring the Primary NPS Server” earlier in this chapter:

- “Obtaining and Installing a Computer Certificate”
- “Configuring NPS Server Properties”
- “Configuring Appropriate Policies”

Next, configure the primary NPS RADIUS server in the first forest with the primary and secondary NPS RADIUS proxies as RADIUS clients. To do this, perform the steps in the “Configuring NPS with RADIUS Clients” subsection of “Configuring the Primary NPS Server” earlier in this chapter. (Instead of the access servers, add the primary and secondary NPS RADIUS proxies as RADIUS clients.)

Configuring the Secondary NPS Server on Another Computer in the First Forest

To configure the secondary NPS RADIUS server on another computer in the first forest, follow the instructions in “Configuring the Secondary NPS Server” earlier in this chapter.

Configuring the Primary NPS Server on a Computer in the Second Forest

To configure the primary NPS RADIUS server on a computer in the second forest, perform the steps in the following subsections of "Configuring the Primary NPS Server" earlier in this chapter on a computer in the second forest:

- "Obtaining and Installing a Computer Certificate"
- "Configuring NPS Server Properties"
- "Configuring Appropriate Policies"

Next, configure the primary NPS RADIUS server in the second forest with the primary and secondary NPS RADIUS proxies as RADIUS clients. To do this, follow the instructions in the "Configuring NPS with RADIUS Clients" subsection of "Configuring the Primary NPS Server" earlier in this chapter (instead of the access servers, add the primary and secondary NPS RADIUS proxies as RADIUS clients).

Configuring the Secondary NPS Server on Another Computer in the Second Forest

To configure the secondary NPS RADIUS server on another computer in the second forest, perform the steps in "Configuring the Secondary NPS Server" earlier in this chapter.

Configuring the Primary NPS RADIUS Proxy

The computer acting as the primary NPS RADIUS proxy is not required to be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server. Because the primary NPS RADIUS proxy computer is not performing authentication or authorization of network access connections, it can be a member of a domain of either forest.

To Configure the Primary NPS RADIUS Proxy for RADIUS Ports and Clients

1. In the Network Policy Server snap-in for the primary NPS RADIUS proxy, configure additional UDP ports for RADIUS messages that are sent by the access servers as needed. By default, NPS uses UDP ports 1812 and 1645 for authentication and UDP ports 1813 and 1646 for accounting.
2. Add the access servers as RADIUS clients by using the instructions in the "Configuring NPS with RADIUS Clients" section of "Configuring the Primary NPS Server" earlier in this chapter.

To Configure the Primary NPS RADIUS Proxy for a Remote RADIUS Server Group Corresponding to the NPS RADIUS Servers in the First Forest

1. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients And Servers.
2. Right-click Remote RADIUS Server Groups, and then click New.

3. In the New Remote RADIUS Server Group dialog box, in the Group Name field, type the group name for the NPS RADIUS servers in the first forest (for example: RADIUS Servers in Forest1). Click Add.
4. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the primary NPS RADIUS server in the first forest. If you specify a name, click Verify to resolve the name to an IP address.
5. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the primary NPS server in the first forest.
6. Click OK to add the server to the list of servers in the group.
7. In the New Remote RADIUS Server Group dialog box, click Add.
8. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the secondary NPS RADIUS server in the first forest.
9. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the secondary NPS server in the first forest.
10. Click OK to add the server to the list of servers in the group, and then click OK again.

To Configure the Primary NPS RADIUS Proxy for a Remote RADIUS Server Group Corresponding to the NPS RADIUS Servers in the Second Forest

1. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients And Servers.
2. Right-click Remote RADIUS Server Groups, and then click New.
3. In the New Remote RADIUS Server Group dialog box, in the Group Name field, type the group name for the NPS RADIUS servers in the second forest (for example: RADIUS Servers in Forest2). Click Add.
4. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the primary NPS RADIUS server in the second forest. If you specify a name, click Verify to resolve the name to an IP address.
5. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the primary NPS RADIUS server in the second forest.
6. Click OK to add the server to the list of servers in the group.
7. In the New Remote RADIUS Server Group dialog box, click Add.
8. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the secondary NPS RADIUS server in the second forest.
9. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the secondary NPS RADIUS server in the second forest.
10. Click OK to add the server to the list of servers in the group, and then click OK again.

To Configure the Primary NPS RADIUS Proxy for a Connection Request Policy to Forward RADIUS Request Messages to the NPS RADIUS Servers in the First Forest

1. In the console tree of the Network Policy Server snap-in, expand Polices, right-click Connection Request Policies, and then click New.
2. On the Specify Connection Request Policy Name And Connection Type page, in the Policy Name box, type the name for the connection request policy (for example: Forward Requests to RADIUS Servers in Forest1). Click Next.
3. On the Specify Conditions page, click Add.
4. In the Select Conditions dialog box, double-click User Name.
5. In the User Name dialog box, type the realm name for all names in the first forest (for example: forest1.example.com), click OK, and then click Next.
6. On the Specify Connection Request Forwarding page, select Forward Requests To The Following Remote RADIUS Server Group For Authentication, and then in the drop-down list, select the remote RADIUS server group for the NPS RADIUS servers in the first forest (for example: RADIUS Servers in Forest1). Click Next.
7. On the Configure Settings page, click Next,
8. On the Completing Connection Request Policy Wizard page, click Finish.

To Configure the Primary NPS RADIUS Proxy for a Connection Request Policy to Forward RADIUS Request Messages to the NPS RADIUS Servers in the Second Forest

1. In the console tree of the Network Policy Server snap-in, expand Policies, right-click Connection Request Policies, and then click New.
2. On the Specify Connection Request Policy Name And Connection Type page, in the Policy Name box, type the name for the connection request policy (for example: Forward Requests to RADIUS Servers in Forest2). Click Next.
3. On the Specify Conditions page, click Add.
4. In the Select Conditions dialog box, double-click User Name.
5. In the User Name dialog box, type the realm name for all names in the second forest (for example: forest2.example.com), click OK, and then click Next.
6. On the Specify Connection Request Forwarding page, select Forward Requests To The Following Remote RADIUS Server Group For Authentication, and then, in the drop-down list, select the remote RADIUS server group for the NPS RADIUS servers in the second forest (for example: RADIUS Servers in Forest2). Click Next.
7. On the Configure Settings page, click Next,
8. On the Completing Connection Request Policy Wizard page, click Finish.

Configuring the Secondary NPS RADIUS Proxy

The computer acting as the secondary NPS RADIUS proxy is not required to be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server. Like the primary NPS RADIUS proxy, the secondary NPS RADIUS proxy computer can be a member of a domain of either forest because it is not performing authentication or authorization of network access connections.

To Configure the Secondary NPS RADIUS Proxy on Another Computer

1. On the primary NPS RADIUS proxy computer, type **netsh nps export *path\file* exportpsk=yes** at a command prompt.

   This command stores the configuration settings, including RADIUS shared secrets, in a text file. The path can be relative, absolute, or a network path.
2. Copy the file created in step 1 to the secondary NPS RADIUS proxy.
3. On the secondary NPS RADIUS proxy computer, type **netsh nps import *path\file*** at a command prompt.

   This command imports all the settings configured on the primary NPS RADIUS proxy into the secondary NPS RADIUS proxy.

Based on the default load-balancing settings of the RADIUS servers in the two remote RADIUS server groups, each NPS RADIUS proxy will distribute the authentication request load equally to the two NPS servers in each forest.

Configuring RADIUS Authentication on the Access Servers

Configure the RADIUS client on your access servers with the following settings:

- The IP address or name of a primary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings.
- The IP address or name of a secondary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings.

To balance the load of RADIUS traffic between the primary and secondary NPS RADIUS proxies, configure half of the access servers with the primary NPS RADIUS proxy as their primary RADIUS server and the secondary NPS RADIUS proxy as their secondary RADIUS server. Configure the other half of the access servers with the secondary NPS RADIUS proxy as their primary RADIUS server and the primary NPS RADIUS proxy as their secondary RADIUS server.

Using RADIUS Proxies to Scale Authentications

When performing authentication for a large number of access clients by using certificate-based authentication or for a large NAP deployment, the volume of RADIUS authentication traffic necessary to keep access clients connected can be substantial. In a large deployment, it is best to spread the load of authentication traffic among multiple NPS server computers. Because you cannot rely on the access servers to consistently or adequately spread their authentication traffic among multiple RADIUS servers, intermediate NPS RADIUS proxies can provide this function.

Without the RADIUS proxies, each access server sends its RADIUS requests to one or multiple RADIUS servers and detects unavailable RADIUS servers. The access server might or might not be balancing the load of RADIUS traffic across multiple RADIUS servers. By using NPS RADIUS proxies, consistent load balancing spreads the load of authentication, authorization, and accounting traffic across all the NPS servers in the organization. Additionally, there is a consistent scheme for failure detection and RADIUS server failover (the detection of an unavailable RADIUS server and avoidance of its use for future authentication requests) and failback (the detection that a previously unavailable RADIUS server is available).

The following configuration is for an organization that uses the following:

- **Active Directory domains** Active Directory domains contain the user accounts, passwords, and dial-in properties that each NPS server requires to authenticate user credentials and evaluate authorization.
- **Multiple NPS servers** To balance a large load of RADIUS authentication, authorization, and accounting traffic, there are multiple NPS servers.
- **Network policies** Network policies are configured to authenticate and authorize network access based on group membership.
- **Two NPS RADIUS proxies** Two NPS RADIUS proxies provide fault tolerance for RADIUS requests that are sent from the access servers.

Figure 9-5 shows the use of NPS RADIUS proxies to balance the load of RADIUS traffic from access servers across multiple NPS servers.

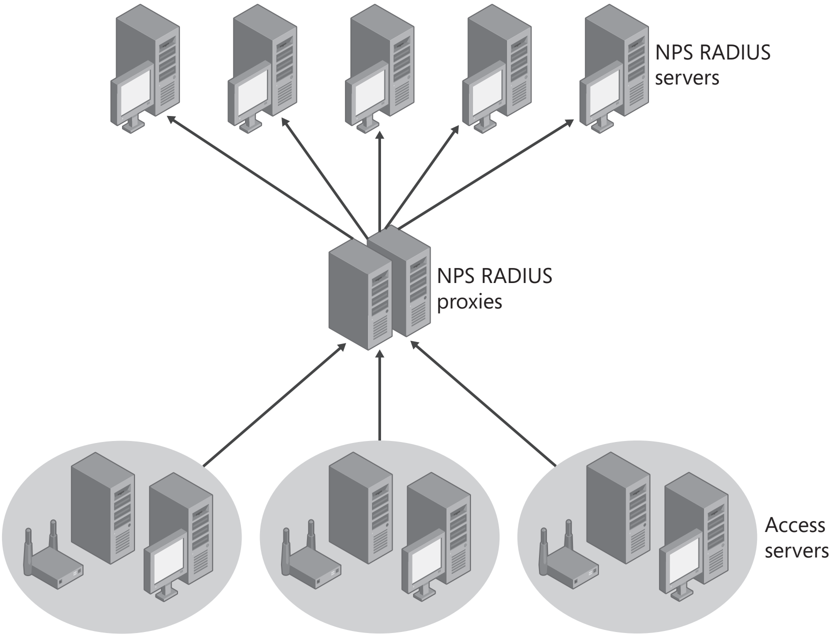

Figure 9-5 Using NPS RADIUS proxies to load-balance RADIUS traffic

To deploy this configuration, do the following:

1. Configure the certificate infrastructure.
2. Configure Active Directory for accounts and groups.
3. Configure NPS as a RADIUS server on multiple computers.
4. Configure the primary NPS RADIUS proxy.
5. Configure the secondary NPS RADIUS proxy.
6. Configure RADIUS authentication and accounting on access servers.

Configuring the Certificate Infrastructure

Follow the instructions in the “Deploying PKI” subsection of “Deployment Steps” earlier in this chapter.

Configuring Active Directory for Accounts and Groups

Follow the instructions in the “Deploying Active Directory” subsection of “Deployment Steps” earlier in this chapter.

Configuring NPS as a RADIUS Server on Multiple Computers

To configure NPS on each NPS server computer, perform on each NPS server computer the steps described in the following subsections of “Configuring the Primary NPS Server” earlier in this chapter:

- “Obtaining and Installing a Computer Certificate”
- “Configuring NPS Server Properties”
- “Configuring Appropriate Policies”

Next, configure each NPS server computer with the primary and secondary NPS RADIUS proxies as RADIUS clients. To do this, perform the steps in the “Configuring NPS with RADIUS Clients” subsection of “Configuring the Primary NPS Server” earlier in this chapter. (Instead of the access servers, add the primary and secondary NPS RADIUS proxies as RADIUS clients.)

> **Note** You can configure each NPS RADIUS server separately rather than configuring an initial NPS RADIUS server and copying its configuration to other NPS RADIUS server computers. This is done so that different RADIUS shared secrets can be used between the NPS RADIUS proxies and the NPS RADIUS server.

Configuring the Primary NPS RADIUS Proxy

The computer acting as the primary NPS RADIUS proxy need not be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server.

To Configure the Primary NPS RADIUS Proxy

1. In the Network Policy Server snap-in, configure additional UDP ports for RADIUS messages that are sent by the access servers if needed.

   By default, NPS uses UDP ports 1812 and 1645 for authentication and UDP ports 1813 and 1646 for accounting.
2. Add the access servers as RADIUS clients of the NPS RADIUS proxy by following the steps in the "Configuring NPS with RADIUS Clients" subsection of "Configuring the Primary NPS Server" earlier in this chapter.
3. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients and Servers.
4. Right-click Remote RADIUS Server Groups, and then click New.
5. In the New Remote RADIUS Server Group box, type the group name for all of the NPS RADIUS servers (for example: RADIUS Servers in the contoso.com Domain).
6. Click Add.
7. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of an NPS RADIUS server. If you specify a name, click Verify to resolve the name to an IP address.
8. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the NPS RADIUS server.
9. Click OK to add the server to the list of servers in the group.
10. Repeat steps 6–9 for each NPS RADIUS server, and then click OK.
11. In the console tree of the Network Policy Server snap-in, expand Policies, right-click Connection Request Policies, and then click New.
12. On the Specify Connection Request Policy Name And Connection Type page, in the Policy Name box, type the name for the connection request policy (for example: Forward Requests to RADIUS Servers in the contoso.com Domain). Click Next.
13. On the Specify Conditions page, click Add.
14. In the Select Conditions dialog box, double-click User Name.
15. In the User Name dialog box, type the realm name for all names in the second forest (for example: forest2.example.com), click OK, and then click Next.
16. On the Specify Connection Request Forwarding page, select Forward Requests To The Following Remote RADIUS Server Group For Authentication, and then in the drop-down list, select the remote RADIUS server group for all of the NPS RADIUS servers in the domain. Click Next.

17. On the Configure Settings page, click Next,
18. On the Completing Connection Request Policy Wizard page, click Finish.

Configuring the Secondary NPS RADIUS Proxy

The computer acting as the secondary NPS RADIUS proxy need not be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server.

To Configure the Secondary NPS RADIUS Proxy on Another Computer

1. On the primary NPS RADIUS proxy computer, type **netsh nps export** ***path\file*** **exportpsk=yes** at a command prompt.

   This command stores the configuration settings, including RADIUS shared secrets, in a text file. The path can be relative, absolute, or a network path.
2. Copy the file created in step 1 to the secondary NPS RADIUS proxy computer.
3. On the secondary NPS RADIUS proxy computer, type **netsh nps import** ***path\file*** at a command prompt. This command imports all the settings configured on the primary NPS RADIUS proxy into the secondary NPS RADIUS proxy.

Based on the default load-balancing settings of the RADIUS servers in the remote RADIUS server group, each NPS RADIUS proxy distributes the authentication request load equally to all of the NPS RADIUS servers.

Configuring RADIUS Authentication on the Access Servers

Configure the RADIUS client on your access servers with the following settings:

- The IP address or name of a primary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings
- The IP address or name of a secondary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings

To balance the load of RADIUS traffic between the primary and secondary NPS RADIUS proxies, configure half of the access servers with the primary NPS RADIUS proxy as their primary RADIUS server and the secondary NPS RADIUS proxy as their secondary RADIUS server. Configure the other half of the access servers with the secondary NPS RADIUS proxy as their primary RADIUS server and the primary NPS RADIUS proxy as their secondary RADIUS server.

Ongoing Maintenance

This section describes the ongoing maintenance for the following components of a Windows authentication infrastructure for network access:

- Active Directory
- PKI

- Group Policy
- RADIUS

Active Directory

It is beyond the scope of this book to describe the ongoing maintenance of an Active Directory infrastructure for an organization of an arbitrary size. For detailed information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/ad*.

The elements of maintaining Active Directory to best support a Windows-based authentication infrastructure for network access are as follows:

- When adding user or computer accounts, ensure that the new accounts have the appropriate security group membership to allow network access. For example, if wireless access is being granted through membership in the WirelessUsers group, add new user or computer accounts to this group or to a group that is a member of this group.
- When adding new domains or forests, ensure that the appropriate trust relationships are created to allow NPS RADIUS servers to authenticate account credentials. Additionally, add the computer accounts of the NPS RADIUS servers to the RAS and IAS Servers security groups of the new domains. If the new domains or forests do not have a trust relationship, use NPS RADIUS proxies to provide cross-domain or cross-forest authentication. For more information, see "Using RADIUS Proxies for Cross-Forest Authentication" earlier in this chapter.

PKI

It is beyond the scope of this book to describe the ongoing maintenance of a PKI for an organization of an arbitrary size. For detailed information, see Windows Server 2008 Help and Support or the resources at *http://www.microsoft.com/pki*.

Group Policy

It is beyond the scope of this book to describe the ongoing maintenance of Group Policy for an organization of an arbitrary size. For detailed information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/gp*.

The elements of maintaining Group Policy to best support a Windows-based authentication infrastructure for network access are as follows:

- When adding new domains or forests, ensure that the appropriate Group Policy objects are applied to the appropriate Active Directory containers to propagate settings for autoenrollment of certificates or configuration settings.

RADIUS

The following sections describe how to maintain the RADIUS component of the network access infrastructure.

Adding a New NPS RADIUS Server to the RADIUS Infrastructure

When you add a new NPS RADIUS server to the RADIUS infrastructure, you must do the following:

1. Register the new NPS server in its default domain.
2. Register the new NPS server in other domains.
3. If the new NPS server is a secondary RADIUS server, obtain and install a computer certificate if needed, and copy the configuration of the primary RADIUS server to the new NPS server.
4. If the new NPS server is a primary RADIUS server, do the following:
   - ❑ Obtain and install a computer certificate.
   - ❑ Configure NPS server properties.
   - ❑ Configure NPS with RADIUS clients.
   - ❑ Configure NPS with the appropriate network policies.
5. Configure access servers (RADIUS clients) to use the new NPS server.
6. If IPsec is being used to protect RADIUS traffic, update Windows Firewall with Advanced Security connection security rules to include protection for RADIUS traffic to and from the new NPS server.

Instructions for these procedures can be found in the “RADIUS Servers” subsection of “Deployment Steps” earlier in this chapter.

Removing an NPS RADIUS Server from the RADIUS Infrastructure

When you remove an NPS RADIUS server from the RADIUS infrastructure, you must do the following:

1. Reconfigure your access servers to remove references to the NPS server that is being removed.
2. Remove the computer account of the NPS server that is being removed from the RAS and IAS Servers security group of its default domain.
3. Remove the computer account of the NPS server that is being removed from the RAS and IAS Servers security group of other domains.
4. If IPsec is being used to protect RADIUS traffic to and from the NPS server that is being removed, update Windows Firewall with Advanced Security connection security rules to remove protection for the NPS server.

Maintaining RADIUS Clients

When you deploy a new access server, such as a new wireless AP for your wireless network, you must do the following:

1. Add the access server as a RADIUS client to either your NPS RADIUS servers or your NPS RADIUS proxies.
2. Configure the access server to use your NPS RADIUS servers or your NPS RADIUS proxies.
3. If IPsec is being used to protect traffic between your RADIUS servers or proxies and the access server, update Windows Firewall with Advanced Security connection security rules to include protection for RADIUS traffic to and from the new access server.

When you remove an access server, you must do the following:

1. Delete the access server as a RADIUS client on either your NPS RADIUS servers or your NPS RADIUS proxies.
2. If IPsec is being used to protect traffic between your RADIUS servers and the access server, update Windows Firewall with Advanced Security connection security rules to remove protection for RADIUS traffic between the access server and the NPS RADIUS servers or proxies.

Troubleshooting Tools

This section describes the troubleshooting tools or the resources that describe troubleshooting tools for the following components of a Windows authentication infrastructure for network access:

- Active Directory
- PKI
- Group Policy
- RADIUS

Active Directory

It is beyond the scope of this book to describe in detail the troubleshooting tools for Active Directory. For additional information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/ad*.

Active Directory–specific troubleshooting issues are described as needed in subsequent chapters to troubleshoot network access or NAP.

PKI

It is beyond the scope of this book to describe in detail the troubleshooting tools for a Windows-based PKI. For additional information, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/pki*.

Digital certificate and PKI-specific troubleshooting issues are described as needed in subsequent chapters to troubleshoot network access or NAP.

Group Policy

It is beyond the scope of this book to describe in detail the troubleshooting tools for Group Policy. For additional information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* by Derek Melber, Group Policy MVP, with the Windows Group Policy Team (Microsoft Press, 2008) Windows Server 2008 Help and Support, or the resources at *http://www.microsoft.com/gp*.

Group Policy–specific troubleshooting issues are described as needed in subsequent chapters to troubleshoot network access or NAP.

RADIUS

To help you gather information to troubleshoot problems with NPS, Microsoft provides the following troubleshooting tools:

- NPS event logging and Windows Event Viewer
- Network Monitor 3.1
- Performance Monitor counters
- SNMP Service

NPS Event Logging and Windows Event Viewer

Use Event Viewer, available from the Administrative Tools program group, to obtain information about hardware and software problems and to monitor all security events, including informational, warning, and error events.

To troubleshoot NPS authentication attempts, view the NPS events in Windows Logs\Security. Viewing the authentication attempts in this log is useful in troubleshooting network policies. When you have multiple network policies configured, you can use the security event log to determine the name of the network policy that either accepted or rejected the connection attempt. Enabling NPS event logging and reading the text of NPS authentication events in the security event log is the most useful tool for troubleshooting failed NPS authentications.

To view the NPS events, configure a filter with the Event Sources option set to Microsoft Windows Security Auditing and the Task Category option set to Network Policy Server.

Both types of logging (rejected authentication requests and successful authentication requests) are enabled by default.

To Configure NPS for Event Logging

1. In the console tree of the Network Policy Server snap-in, right-click NPS, and then click Properties.
2. On the General tab, select each required check box, and then click OK.

Network Monitor 3.1

You can use Network Monitor 3.1 (or later) or a commercial packet analyzer (also known as a *network sniffer*), to capture and view RADIUS authentication and accounting messages that are sent to and from the NPS server. Network Monitor 3.1 includes a RADIUS parser, which you can use to view the attributes of a RADIUS message and troubleshoot network access or NAP issues.

> **On the Disc** You can link to the download site for Network Monitor from the companion CD-ROM.

Reliability and Performance Counters

You can use the Reliability and Performance snap-in to monitor the resource use of specific components and program processes. With Performance Monitor, which is in the Reliability and Performance snap-in, you can use charts and reports to determine how efficiently your server uses NPS and both identify and troubleshoot potential problems.

You can use Performance Monitor to monitor the following NPS-related performance objects:

- NPS Accounting Clients
- NPS Accounting Server
- NPS Authentication Clients
- NPS Authentication Server

SNMP Service

You can use the Simple Network Management Protocol (SNMP) service to monitor status information for your NPS server. NPS supports the RADIUS Authentication Server Management Information Base (MIB), as specified in RFC 2619, and the RADIUS Accounting Server MIB, as specified in RFC 2621.

Chapter Summary

A Windows-based network access infrastructure consists of Active Directory, PKI, Group Policy, and RADIUS components. Active Directory stores user and computer account credentials and properties and provides an infrastructure to deploy centrally configured user and computer configuration Group Policy settings. A PKI issues and validates digital certificates used in different types of network access scenarios or NAP enforcement methods. Group Policy settings can instruct computers to automatically request specific types of certificates or configure network access and protection settings. RADIUS provides a standard protocol and centralized management of network access authorization, authentication, and accounting.

The combination of Active Directory, PKI, Group Policy, and RADIUS creates a Windows-based infrastructure that provides centralized authentication for 802.11 wireless access, 802.1X wired access, dial-up or VPN-based remote access connections, and dial-up or VPN-based site-to-site connections. The combination of PKI, Group Policy, and RADIUS creates a Windows-based infrastructure that provides centralized configuration and validation of system health status for NAP.

Additional Information

For additional information about Active Directory, see the following:

- *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit* (both from Microsoft Press, 2008)
- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Microsoft Windows Server Active Directory (*http://www.microsoft.com/ad*)

For additional information about PKI, see the following:

- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Microsoft Public Key Infrastructure for Windows Server (*http://www.microsoft.com/pki*)
- *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008)

For additional information about Group Policy, see the following:

- *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* (Microsoft Press, 2008)
- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Microsoft Windows Server Group Policy (*http://www.microsoft.com/gp*)

For additional information about RADIUS and NPS, see the following:

- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Network Policy Server (*http://www.microsoft.com/nps*)
- RFC 2548, "Microsoft Vendor-Specific RADIUS Attributes"
- RFC 2619, "RADIUS Authentication Server MIB"
- RFC 2621, "RADIUS Accounting Server MIB"
- RFC 2865, "Remote Authentication Dial-In User Service (RADIUS)"
- RFC 2866, "RADIUS Accounting"
- RFC 2867, "RADIUS Accounting Modifications for Tunnel Protocol Support"
- RFC 2868, "RADIUS Attributes for Tunnel Protocol Support"
- RFC 2869, "RADIUS Extensions"
- RFC 3162, "RADIUS and IPv6"
- RFC 3579, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)"

For additional information about Windows-based network access, see the following:

- Chapter 10, "IEEE 802.11 Wireless Networks"
- Chapter 11, "IEEE 802.1X-Authenticated Wired Networks"
- Chapter 12, "Remote Access VPN Connections"
- Chapter 13, "Site-to-Site VPN Connections"

For additional information about NAP, see the following:

- Chapter 14, "Network Access Protection Overview"
- Chapter 15, "Preparing for Network Access Protection"
- Chapter 16, "IPsec Enforcement"
- Chapter 17, "802.1X Enforcement"
- Chapter 18, "VPN Enforcement"
- Chapter 19, "DHCP Enforcement"
- The Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Microsoft Network Access Protection (*http://www.microsoft.com/nap*)

Chapter 10

IEEE 802.11 Wireless Networks

This chapter provides information about how to design, deploy, maintain, and troubleshoot Institute of Electrical and Electronic Engineers (IEEE) 802.11 wireless networks. Once deployed, the protected wireless network solution can be modified for the 802.1X Enforcement method of Network Access Protection (NAP) as described in Chapter 17, "802.1X Enforcement."

This chapter assumes that you understand the role of Active Directory, public key infrastructure (PKI), Group Policy, and Remote Authentication Dial-In User Service (RADIUS) elements of a Microsoft Windows–based authentication infrastructure for network access. For more information, see Chapter 9, "Authentication Infrastructure."

Concepts

IEEE 802.11 wireless local area network (LAN) networking provides the following benefits:

- Wireless connections can extend or replace a wired infrastructure in situations where it is costly, inconvenient, or impossible to lay cables. This benefit includes the following:
  - ❑ To connect the networks in two buildings separated by a physical, legal, or financial obstacle, you can either use a link provided by a telecommunications vendor (for a fixed installation cost and ongoing recurring costs), or you can create a point-to-point wireless link using wireless LAN technology (for a fixed installation cost but no recurring costs). Eliminating recurring telecommunications charges can provide significant cost savings to organizations.
  - ❑ Wireless LAN technologies can be used to create a temporary network, which is in place for only a specific amount of time. For example, you can set up a wireless network for a convention or trade show rather than deploying the physical cabling required for a traditional Ethernet network.
  - ❑ Some types of buildings, such as historical buildings, might be governed by building codes that prohibit the use of wiring, making wireless networking an important alternative.
- The wiring-free aspect of wireless LAN networking is also attractive to homeowners who want to connect the various computers in their home together without having to drill holes and pull network cables through walls and ceilings.

- Increased productivity for the mobile employee. This benefit includes the following:
  - The mobile user whose primary computer is a laptop or notebook computer can change location and always remain connected to the network. This enables the mobile user to travel to various places–meeting rooms, hallways, lobbies, cafeterias, classrooms, and so forth–and still have access to networked data. Without wireless access, the user has to carry cabling and is restricted to working near a network jack.
  - Wireless LAN networking is well suited for environments where movement is required. For example, retail environments can benefit when employees use a wireless laptop or palmtop computer to enter inventory information directly into the store database from the sales floor.
  - Even if no wireless infrastructure is present, wireless laptop computers can still form their own ad-hoc networks to communicate and share data with each other.
- Easy access to the Internet in public places. Beyond the corporate campus, access to the Internet and even corporate sites can be made available through public wireless "hot spot" networks. Airports, restaurants, rail stations, and common areas throughout cities can be provisioned to provide this service. When the traveling worker reaches his or her destination, perhaps meeting a client at the client's corporate office, limited access can be provided to the traveling worker through the local wireless network. The network can recognize that a user is from another corporation and create a connection that is isolated from the local corporate network but provides Internet access to the visiting user. Wireless infrastructure providers are enabling wireless connectivity in public areas around the world. Many airports, conference centers, and hotels provide wireless access to the Internet for their visitors.

Support for IEEE 802.11 Standards

The Windows Server 2008, Windows Vista, Windows XP, and Windows Server 2003 operating systems provide built-in support for 802.11 wireless LAN networking. An installed 802.11 wireless LAN network adapter appears as a wireless network connection in the Network Connections folder. Although there is built-in support for 802.11 wireless LAN networking, the wireless components of Windows are dependent upon the following:

- **The capabilities of the wireless network adapter** The installed wireless network adapter must support the wireless LAN or wireless security standards that you require. For example, Windows Vista supports configuration options for the Wi-Fi Protected Access (WPA) security standard. However, if the wireless network adapter does not support WPA, you cannot enable or configure WPA security options.
- **The capabilities of the wireless network adapter driver** To allow you to configure wireless network options, the driver for the wireless network adapter must support the reporting of all of its capabilities to Windows. Verify that the driver for your wireless

network adapter was written for the capabilities of Windows Vista or Windows XP and is the most current version by checking Microsoft Update or the Web site of the wireless network adapter vendor.

Table 10-1 lists the IEEE wireless standards supported by Windows and by wireless network adapters, their maximum bit rate, range of frequencies, and their typical usage.

Table 10-1 802.11 Standards

Standard	Maximum Bit Rate	Range of Frequencies	Usage
802.11	2 megabits per second (Mbps)	S-Band Industrial, Scientific, and Medical (ISM) frequency range (2.4 to 2.5 GHz)	Obsolete. Not widely used.
802.11b	11 Mbps	S-Band ISM	Widely used.
802.11a	54 Mbps	C-Band ISM (5.725 to 5.875 GHz)	Not widely used due to expense and limited range.
802.11g	54 Mbps	S-Band ISM	Widely used. 802.11g devices are backward-compatible with 802.11b devices.
802.11n (standards development in progress)	250 Mbps	C-Band and S-Band ISM	Pre-standard ratification devices are available starting in August 2007. 802.11n devices can be backward-compatible with 802.11a, b, and g devices.

Note The S-Band ISM uses the same frequency range as microwave ovens, cordless phones, baby monitors, wireless video cameras, and Bluetooth devices. The C-Band ISM uses the same frequency range as newer cordless phones and other devices. Due to this overlapping use, there might be contention when multiple devices are active at the same time.

802.11 Operating Modes

Wireless LAN networks for all the IEEE 802.11 standards use the following operating modes:

- **Infrastructure mode** The wireless network contains at least one wireless access point (AP), a device that bridges wireless-based computers to each other and to a wired network such as the Internet or an intranet.
- **Ad-hoc mode** The wireless network contains no wireless APs. Wireless-based computers connect and communicate directly with each other. This chapter does not describe ad-hoc mode wireless networks.

Regardless of the operating mode, a *Service Set Identifier* (SSID), also known as the wireless network name, identifies a specific wireless network. You configure the SSID on the wireless AP for infrastructure mode or the initial wireless client for ad-hoc mode. The wireless AP or the initial wireless client periodically advertises the SSID so that other wireless nodes can discover and join the wireless network.

Wireless Security

Although IEEE 802.11 wireless LAN technologies provide the benefits previously described, they introduce security issues that do not exist for wired networks. Unlike the closed cabling system of an Ethernet network, which can be physically secured, wireless frames are sent as radio transmissions that propagate beyond the physical confines of your office. Any computer within range of the wireless network can receive wireless frames and send its own. Without protecting your wireless network, malicious users can use your wireless network to access your private information or launch attacks against your computers or other computers across the Internet.

To protect your wireless network, you must use authentication and encryption, described as follows:

- Authentication requires that computers provide either valid account credentials (such as a user name and password) or proof that they have been configured with a specific authentication key before being allowed to send data frames on the wireless network. Authentication prevents malicious users from being able to join your wireless network.
- Encryption requires that the content of all wireless data frames be encrypted so that only the designated receiver can interpret its contents. Encryption prevents malicious users from capturing wireless frames sent on your wireless network and determining sensitive data. Encryption also helps prevent malicious users from sending valid frames and accessing your private resources or the Internet.

IEEE 802.11 wireless LANs support the following security standards:

- IEEE 802.11
- IEEE 802.1X
- Wi-Fi Protected Access (WPA)
- Wi-Fi Protected Access 2 (WPA2)

IEEE 802.11

The original IEEE 802.11 standard defined the open system and shared key authentication methods for authentication and Wired Equivalent Privacy (WEP) for encryption. WEP can use either 40-bit or 104-bit encryption keys. However, the original IEEE 802.11 security standard has proven to be relatively weak and because there was no specified method for WEP encryption key management, cumbersome for widespread public and private deployment. Because of its susceptibility to attack and the widespread support of newer security standards, such as WPA and WPA2, its use is highly discouraged.

IEEE 802.1X

IEEE 802.1X was a standard that existed for Ethernet switches and was adapted to 802.11 wireless LANs to provide much stronger authentication than the original 802.11 standard. IEEE 802.1X authentication is designed for medium and large wireless LANs that contain an authentication infrastructure consisting of Remote Authentication Dial-In User Service (RADIUS) servers and account databases, such as the Active Directory domain service.

IEEE 802.1X prevents a wireless node from joining a wireless network until the node is successfully authenticated and authorized. Authentication verifies that wireless clients have valid account credentials and prevents users without valid credentials from being able to join your wireless network. Authorization verifies that the wireless client is allowed to make a connection to the wireless AP. IEEE 802.1X uses the Extensible Authentication Protocol (EAP) to exchange authentication credentials. IEEE 802.1X authentication can be based on different EAP authentication methods, such as those using user name and password credentials or a digital certificate.

To address the key management issues of the original 802.11 standard, 802.1X authentication can produce dynamic WEP keys, which are mutually determined by the wireless client and RADIUS server. The RADIUS server sends the WEP key to the wireless AP after authentication completes. The combination of WEP encryption and dynamic keys determined for each 802.1X authentication is known as dynamic WEP.

WPA

Although 802.1X addresses the weak authentication and key management issues of the original 802.11 standard, it provides no solution to the weaknesses of the WEP encryption algorithm. While the IEEE 802.11i wireless LAN security standard, which will be discussed in the section titled “WPA2” later in this chapter, was being finalized, the Wi-Fi Alliance, an organization of wireless equipment vendors, created an interim standard known as Wi-Fi Protected Access (WPA). WPA replaces WEP with a much stronger encryption method known as the Temporal Key Integrity Protocol (TKIP). WPA also allows the optional use of the Advanced Encryption Standard (AES) for encryption.

WPA is available in two different modes:

- **WPA-Enterprise** Uses 802.1X authentication and is designed for medium and large infrastructure mode networks
- **WPA-Personal** Uses a preshared key (PSK) for authentication and is designed for small office/home office (SOHO) infrastructure mode networks

WPA2

The IEEE 802.11i standard formally replaces WEP and the other security features of the original IEEE 802.11 standard. Wi-Fi Protected Access 2 (WPA2) is a product certification available through the Wi-Fi Alliance that certifies wireless equipment as being compatible with the IEEE 802.11i standard. The goal of WPA2 certification is to support the additional mandatory security features of the IEEE 802.11i standard that are not already included for products that support WPA. For example, WPA2 requires support for both TKIP and AES encryption. WPA2 includes fast roaming techniques, such as Pairwise Master Key (PMK) caching and pre-authentication.

How It Works: Fast Roaming for WPA2

When a wireless client authenticates using 802.1X, there are a series of messages sent between the wireless client and the wireless AP to exchange credentials (802.1X authentication) and to determine the pairwise transient keys (the 4-way handshake). The pairwise transient keys are used for encryption and data integrity of WPA2-protected wireless data frames. This message exchange introduces a delay in the connection process. When a wireless client roams from one wireless AP to another, the delay to perform 802.1X authentication can cause noticeable interruptions in network connectivity, especially for time-dependent traffic, such as voice-based or video-based data streams. To minimize the delay associated with roaming to another wireless AP, WPA2 wireless equipment can optionally support PMK caching and preauthentication.

PMK Caching

As a wireless client roams from one wireless AP to another, it must perform a full 802.1X authentication with each wireless AP. WPA2 allows the wireless client and the wireless AP to cache the results of a full 802.1X authentication so that if a client roams back to a wireless AP with which it has previously authenticated, the wireless client needs to perform only the 4-way handshake and determine new pairwise transient keys. In the Association Request frame, the wireless client includes a PMK identifier that was determined during the initial authentication and stored with both the wireless client and wireless AP's PMK cache entries. PMK cache entries are stored for a finite amount of time as configured on the wireless client and the wireless AP.

To make the transition faster for wireless networking infrastructures that use a switch that acts as the 802.1X authenticator, Windows Server 2008 and Windows Vista calculate the PMK identifier value so that the PMK as determined by the 802.1X authentication with the switch can be reused when roaming between wireless APs that are attached to the same switch. This practice is known as *opportunistic PMK caching*.

Preauthentication

With preauthentication, a WPA2 wireless client can optionally perform 802.1X authentications with other wireless APs within its range while connected to its current wireless

AP. The wireless client sends preauthentication traffic to the additional wireless AP over its existing wireless connection. After preauthenticating with a wireless AP and storing the PMK and its associated information in the PMK cache, a wireless client that connects to a wireless AP with which it has preauthenticated needs to perform only the 4-way handshake.

WPA2 clients that support preauthentication can preauthenticate only with wireless APs that advertise their preauthentication capability in Beacon and Probe Response frames.

WPA2 is available in two different modes:

- **WPA2-Enterprise** Uses 802.1X authentication and is designed for medium and large infrastructure mode networks
- **WPA2-Personal** Uses a PSK for authentication and is designed for SOHO infrastructure mode networks

Table 10-2 summarizes the 802.11 wireless LAN security standards.

Table 10-2 802.11 Wireless LAN Security Standards

Security Standard	Authentication Methods	Encryption Methods	Encryption Key Size (bits)	Comments
IEEE 802.11	Open system and shared key	WEP	40 and 104	Weak authentication and encryption. Use is highly discouraged.
IEEE 802.1X	EAP authentication methods	N/A	N/A	Strong EAP methods provide strong authentication.
WPA-Enterprise	802.1X	TKIP and AES (optional)	128	Strong authentication (with strong EAP method) and strong (TKIP) or very strong (AES) encryption.
WPA-Personal	PSK	TKIP and AES (optional)	128	Strong authentication (with strong PSK) and strong (TKIP) or very strong (AES) encryption.
WPA2-Enterprise	802.1X	TKIP and AES	128	Strong authentication (with strong EAP method) and strong (TKIP) or very strong (AES) encryption.
WPA2-Personal	PSK	TKIP and AES	128	Strong authentication (with strong PSK) and strong (TKIP) or very strong (AES) encryption.

Windows Server 2008 and Windows Vista support the following security standards for 802.11 wireless LAN networking (the wireless network adapter and driver must also support the standard):

- 802.11 with WEP
- 802.1X
- WPA-Enterprise
- WPA-Personal
- WPA2-Enterprise
- WPA2-Personal

Note Unless stated otherwise, all subsequent references to WPA2 refer to WPA2-Enterprise, and references to WPA refer to WPA-Enterprise.

Components of 802.11 Wireless Networks

Figure 10-1 shows the components of Windows-based 802.11 protected wireless networks.

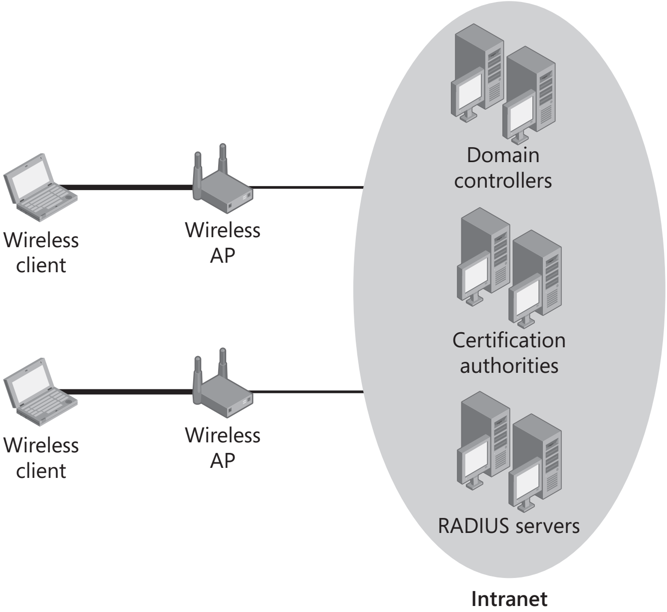

Figure 10-1 Components of Windows-based 802.11 protected wireless networks

The components are:

- **Wireless clients** Initiate wireless connections to wireless APs and communicate with intranet resources or other wireless clients once connected
- **Wireless APs** Listen for wireless connection attempts, enforce authentication and connection requirements, and forward frames between wireless clients and intranet resources
- **RADIUS servers** Provide centralized authentication and authorization processing and accounting for network access attempts from wireless APs and other types of access servers
- **Active Directory domain controllers** Validate user credentials for authentication and provide account information to the RADIUS servers to evaluate authorization
- **Certification authorities** Part of the PKI that issues computer or user certificates to wireless clients and computer certificates to RADIUS servers

Planning and Design Considerations

When deploying a protected 802.11 wireless network solution, you must consider the following for planning and design issues:

- Wireless security technologies
- Wireless authentication modes
- Intranet infrastructure
- Wireless AP placement
- Authentication infrastructure
- Wireless clients
- PKI
- 802.1X Enforcement with NAP

Wireless Security Technologies

Wireless security technologies are a combination of a wireless security standard (WPA2 or WPA) and an EAP authentication method. To authenticate the computer or the user that is attempting to make a protected wireless connection, Windows Server 2008 and Windows Vista support the following EAP authentication methods:

- EAP-TLS
- Protected EAP (PEAP)-TLS
- PEAP-Microsoft Challenge Handshake Authentication Protocol version 2 (PEAP-MS-CHAP v2)

EAP-TLS and PEAP-TLS are used in conjunction with a PKI and computer certificates, user certificates, or smart cards. With EAP-TLS, the wireless client sends its computer certificate, user certificate, or smart card certificate for authentication, and the RADIUS server sends its computer certificate for authentication. By default, the wireless client validates the RADIUS server's certificate. With PEAP-TLS, the wireless client and RADIUS server create an encrypted TLS session, and then the wireless client and RADIUS server exchange certificates. PEAP-TLS is the strongest authentication method because the certificate exchange between the wireless client and the RADIUS server is encrypted.

In the absence of computer certificates, user certificates, or smart cards, use PEAP-MS-CHAP v2. PEAP-MS-CHAP v2 is a password-based authentication method in which the exchange of authentication messages is protected with an encrypted TLS session, making it much more difficult for a malicious user to determine the password of a captured authentication exchange with an offline dictionary attack.

Despite the encrypted TLS session, however, both EAP-TLS and PEAP-TLS are much stronger than PEAP-MS-CHAP v2 because they do not rely on passwords.

Design Choices for Wireless Security Technologies

Microsoft recommends that you use one of the following combinations of wireless security technologies (in order of most to least secure):

- WPA2 with AES encryption, PEAP-TLS or EAP-TLS authentication, and both user and computer certificates
- WPA2 with AES encryption, PEAP-MS-CHAP v2 authentication, and a requirement for users to create strong user passwords
- WPA with EAP-TLS or PEAP-TLS authentication and both user and computer certificates
- WPA with PEAP-MS-CHAP v2 authentication and a requirement for users to create strong user passwords

Requirements for Wireless Security Technologies

The requirements for wireless security technologies are the following:

- For a protected wireless network, you must use either WPA or WPA2. If you use WEP, even dynamic WEP, your wireless network will not be secure. Dynamic WEP should not be used except temporarily when transitioning to a WPA2 or WPA-based security configuration.
- EAP-TLS or PEAP-TLS requires the installation of a computer certificate on the RADIUS server and a computer certificate, user certificate, or smart card on all wireless client computers. To validate the RADIUS servers' computer certificates, the root certification authority (CA) certificate of the issuing CA of the RADIUS server computer certificates must be installed on all wireless client computers. To validate the wireless clients'

computer or user certificates, the root CA certificate of the issuing CA of the wireless client certificates must be installed on each of the RADIUS servers.

- PEAP-MS-CHAP v2 requires the installation of computer certificates on each of the RADIUS servers. It also requires that the root CA certificates of the RADIUS server computer certificates be installed on each of the wireless client computers.
- For WPA2, some wireless equipment might need to be replaced. Older wireless equipment that supports only 802.11 can typically be upgraded to support WPA but not WPA2.
- If you are planning to eventually deploy the 802.1X Enforcement method of NAP, you should use a PEAP-based authentication method, such as PEAP-MS-CHAP v2 or PEAP-TLS.

Best Practices for Wireless Security Technologies

The best practices for wireless security technologies are as follows:

- Do not configure your wireless APs to use SSID suppression. The SSID (also known as the wireless network name) is by default included in the Beacon frames sent by wireless APs. Configuring your wireless APs to suppress the advertising of the SSID information element in Beacon frames does prevent the casual wireless client from discovering your wireless network. However, SSID suppression does not prevent a more sophisticated hacker from capturing other types of wireless management frames sent by your wireless AP and determining your SSID. Wireless networks with SSID suppression enabled are known as *non-broadcast* or *hidden* networks.

  Besides being a weak form of wireless network name privacy, non-broadcast wireless networks also create problems for authorized wireless clients that want to automatically connect to the non-broadcast wireless network. For example, because the wireless network name is not being advertised, the wireless client must send Probe-Request messages containing the wireless network name in an attempt to locate a wireless AP for the wireless network. These messages advertise the name of the wireless network, reducing the privacy of the wireless configuration of the wireless client.
- Do not use media access control (MAC) address filtering. MAC address filtering allows you to configure your wireless APs with the set of MAC addresses for allowed wireless clients. MAC address filtering adds administrative overhead in order to keep the list of allowed MAC addresses current and does not prevent a hacker from spoofing an allowed MAC address.
- If you must use PEAP-MS-CHAP v2, require the use of strong passwords on your network. Strong passwords are long (longer than 8 characters) and contain a mixture of upper and lower case letters, numbers, and punctuation. In an Active Directory domain, use Group Policy settings in Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy to enforce strong user passwords requirements.

Wireless Authentication Modes

Windows-based wireless clients can perform authentication using the following modes:

- **Computer-only** Windows performs 802.1X authentication with computer credentials before displaying the Windows logon screen. This allows the wireless client to have access to networking resources, such as Active Directory domain controllers, before the user logs on. Windows does not attempt authentication with user credentials after the user logs on.
- **User-only** By default, Windows performs 802.1X authentication with user credentials after the user logon process has completed. Windows does not attempt authentication with computer credentials before the user logon.
- **Computer-or-user** Windows performs an 802.1X authentication with computer credentials before displaying the Windows logon screen. Windows performs another 802.1X authentication with user credentials either after the user has logged on or when the wireless client roams to a new wireless AP.

Problems with the default behavior of user-only authentication mode are as follows:

- A user cannot perform an initial domain logon to a computer because locally cached credentials for the user's user account are not available and there is no connectivity to the domain controller to authenticate new logon credentials.
- Domain logon operations will not be successful because there is no connectivity to the domain controllers of the Active Directory domain during the user logon process. Logon scripts, Group Policy updates, and user profile updates will fail, resulting in Windows event log errors.

Some network infrastructures use different virtual LANs (VLANs) to separate wireless clients that have authenticated with computer credentials from wireless clients that have authenticated with user credentials. If the user-level authentication to the wireless network and the switch to the user-authenticated VLAN occurs after the user logon process, a Windows wireless client will not have access to resources on the user-authenticated VLAN—such as Active Directory domain controllers—during the user logon process. This can lead to unsuccessful initial logons and domain logon operations such as logon scripts, Group Policy updates, and user profile updates.

To address the availability of network connectivity when performing user logon in user-only authentication mode and user-or-computer authentication mode when using separate VLANs, Windows Server 2008 and Windows Vista wireless clients support Single Sign On. With Single Sign On, you can specify that wireless network authentication with user credentials occur before the user logon process. To enable and configure Single Sign On, you can use the Wireless Network (IEEE 802.11) Policies Group Policy extension to configure a Windows Vista policy, or you can run **netsh wlan** with the appropriate parameters. For more information, see the section titled "Configuring Wireless Clients" later in this chapter.

Requirements for Wireless Authentication Modes

Only wireless clients running Windows Server 2008 or Windows Vista support Single Sign On.

Best Practices for Wireless Authentication Modes

Best practices for wireless authentication modes are as follows:

- Use user-or-computer authentication mode; user authentication occurs after user logon. This is the default authentication mode.
- If you are using user-only authentication mode, configure your wireless profiles to enable Single Sign On and perform wireless authentication with user credentials before user logon to prevent initial and domain logon problems.
- If you are using different VLANs for computer-authenticated and user-authenticated wireless clients and computer-or-user authentication mode, configure your wireless profiles to enable Single Sign On and perform wireless authentication with user credentials before user logon to prevent initial and domain logon problems.

Intranet Infrastructure

Wireless clients need the same Transmission Control Protocol/Internet Protocol (TCP/IP) configuration settings and connectivity as wired clients, but there are differences in how you should configure wireless clients because of their inherent mobility. For this reason, place your wireless clients on different subnets than your wired clients rather than have a mixture of wired and wireless clients on the same subnet.

Subnet Design for Wireless Clients

Creating separate subnets for your wireless clients provides the following benefits:

- Wired network components do not need to draw from the same pool of existing IPv4 addresses as your wireless clients.
- Wireless clients are easier to identify from their IPv4 and IPv6 address prefixes, which makes it easier to manage and troubleshoot wireless clients.
- Separate IPv4 subnets give you increased control over DHCP lease times.
- You can associate each of your physical subnets (both wireless and wired) with sites within Active Directory, which allows you to assign Group Policy settings to specific subnets.
- If all of your wireless APs are on the same subnet, your wireless clients can seamlessly perform network-layer roaming.

Network-layer roaming occurs when a wireless client connects to a different wireless AP for the same wireless network within the same subnet. For network-layer roaming, the wireless client renews its current DHCP configuration. When a wireless client connects to a different wireless AP for the same wireless network that is on a different subnet, the wireless client gets a new DHCP configuration that is relevant to that new subnet. When you cross a subnet boundary, applications that cannot handle a change of IPv4 or IPv6 address, such as some e-mail applications, might fail.

When creating an IPv4 subnet prefix for your wireless clients, consider that you need at least one IPv4 address for the following:

- Each wireless AP's LAN interface that is connected to the wireless subnet.
- Each router interface that is connected to the wireless subnet.
- Any other TCP/IP-capable host or device that is attached to the wireless subnet.
- Each wireless client that can connect to the wireless network. If you underestimate this number, Windows wireless clients that connect after all of the available IPv4 addresses have been assigned through DHCP to connected wireless clients will automatically configure an IP address with no default gateway using Automatic Private IP Addressing (APIPA). This configuration does not allow connectivity to the intranet. Wireless clients with APIPA configurations will periodically attempt to obtain a DHCP configuration.

Because each IPv6 subnet can support a very large number of hosts, you do not need to determine the number of IPv6 addresses needed for the IPv6 subnet prefix.

DHCP Design for Wireless Clients

With different subnets for wired and wireless clients, you must configure separate DHCP scopes. Because wireless clients can easily roam from one wireless subnet to another, you should configure the lease for the DHCP scopes to have a shorter duration for wireless subnets than for wired subnets.

The typical lease duration for a DHCP scope for wired networks is a specified number of days. Because wireless clients do not release their addresses when roaming to a new subnet, you should shorten the lease duration to several hours for DHCP scopes corresponding to wireless subnets. By setting a shorter lease duration for wireless subnets, the DHCP server will automatically make IPv4 addresses that are no longer being used by wireless clients available for reuse throughout the day instead of leaving the addresses unavailable for days. When determining the optimal lease duration for the wireless clients in your environment, keep in mind the additional processing load that the shorter lease duration places on your DHCP server.

For more information about configuring DHCP scopes, see Chapter 3, "Dynamic Host Configuration Protocol."

Wireless AP Placement

An important and time-consuming task in deploying a wireless LAN is determining where to place the wireless APs in your organization. Wireless APs must be placed to provide seamless coverage across the floor, building, or campus. With seamless coverage, wireless users can roam from one location to another without experiencing an interruption in network connectivity, except for a change in IPv4 and IPv6 addresses when crossing a subnet boundary. Determining where to place your wireless APs is not as simple as installing them and turning them on. Wireless LAN technologies are based on propagation of a radio signal, which can be obstructed, reflected, shielded, and interfered with.

When planning the deployment of wireless APs in an organization, you should take the following design elements into consideration (as described in the following sections):

- Wireless AP requirements
- Channel separation
- Signal propagation modifiers
- Sources of interference
- Number of wireless APs

> **Note** For additional specifications and guidelines for placing wireless APs, see the manufacturer's documentation for the wireless APs and the antennas used with them.

Wireless AP Requirements

You must identify the requirements for your wireless APs, which might include the following features:

- WPA
- WPA2
- 802.1X and RADIUS
- 802.11a, b, g, and n

Depending on your budget and bandwidth requirements, you might need wireless APs that support 802.11b, 802.11a, 802.11g, 802.11n, or a combination of technologies.

- **Building or fire code compliance** The plenum area (the space between the suspended ceiling and the ceiling) is regulated by building and fire codes. Therefore, for plenum placement of APs and associated wiring, you must purchase wireless APs that are fire-rated and in compliance with building and fire codes. If you place your wireless APs in the plenum area, you must determine the best method for powering the wireless

APs. Consult with the wireless AP manufacturer to determine how to meet the power requirements for the wireless APs. Some wireless APs can receive electrical power through the Ethernet cable that connects them to the wired network.

- **Preconfiguration and remote configuration** Preconfiguring the wireless APs before installing them on location can speed up the deployment process and can save labor costs because less-skilled workers can perform the physical installation. You can preconfigure wireless APs by using the console port (serial port), Telnet, or a Web server that is integrated with the wireless AP. Regardless of whether you decide to preconfigure the wireless APs, make sure that you can access them remotely, configure the wireless APs remotely through a vendor-supplied configuration tool, or upgrade the wireless APs by using scripts.
- **Antenna types** Verify that the wireless AP supports different types of antennas. For example, in a building with multiple floors, a loop antenna—which propagates the signal equally in all directions except vertically—might work best.

> **Note** For information about which type of antenna will work best for your wireless WLAN deployment, see the documentation for your wireless APs.

- **IPsec support** Although not a requirement, if possible, choose wireless APs that use Internet Protocol security (IPsec) and Encapsulating Security Payload (ESP) with encryption to provide data confidentiality for RADIUS traffic sent between wireless APs and RADIUS servers. Use Triple Data Encryption Standard (3DES) encryption and, if possible, certificates for Internet Key Exchange (IKE) main mode authentication.

Channel Separation

Direct communication between an 802.11b or 802.11g wireless network adapter and a wireless AP occurs over a common channel, which corresponds to a frequency range in the S-Band ISM. You configure the wireless AP for a specific channel, and the wireless network adapter automatically configures itself to the channel of the wireless AP with the strongest signal.

To reduce interference between 802.11b wireless APs, ensure that wireless APs with overlapping coverage volumes use unique frequency channels. The 802.11b or 802.11g standards reserve 14 channels for use with wireless APs. Within the United States, the Federal Communications Commission (FCC) allows channels 1 through 11. In most of Europe, you can use channels 1 through 13. In Japan, you have only one choice: channel 14. Figure 10-2 shows the channel overlap for 802.11b and 802.11g wireless APs in the United States.

To prevent signals from adjacent wireless APs from interfering with one another, you must set their channel numbers so that they are at least five channels apart. To get the most usable channels in the United States, you can set your wireless APs to use one of three channels: 1, 6, or 11. If you need fewer than three usable channels, ensure that the channels you choose maintain the five-channel separation.

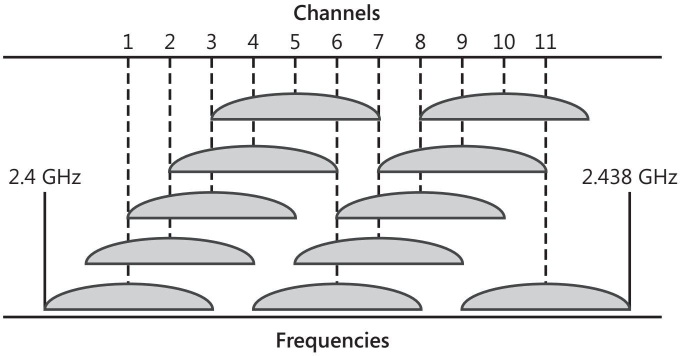

Figure 10-2 Channel overlap for 802.11b and 802.11g wireless APs in the United States

Figure 10-3 shows an example of a set of wireless APs deployed in multiple floors of a building so that overlapping signals from adjacent wireless APs use different usable channel numbers.

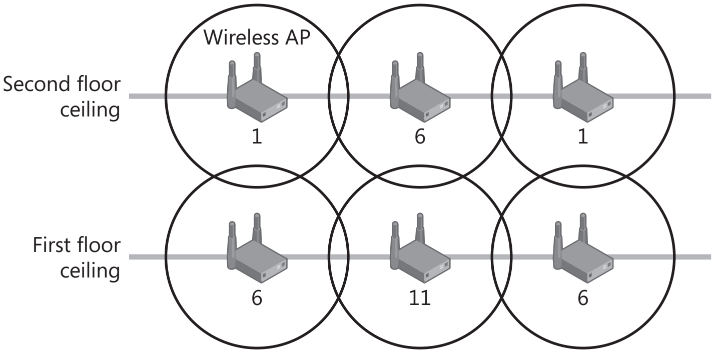

Figure 10-3 Example of assigning 802.11b channel numbers

Signal Propagation Modifiers

The wireless AP is a radio transmitter and receiver that has a limited range. The volume around the wireless AP for which you can send and receive wireless data for any of the supported bit rates is known as the *coverage volume*. (Many wireless references use the term *coverage area*; however, wireless signals propagate in three dimensions.) The shape of the coverage volume depends on the type of antenna used by the wireless AP and the presence of signal propagation modifiers and other interference sources.

With an idealized omnidirectional antenna, the coverage volume is a series of concentric spherical shells of signal strengths corresponding to the different supported bit rates. Figure 10-4 shows an example of the idealized coverage volume for 802.11b and an omnidirectional antenna.

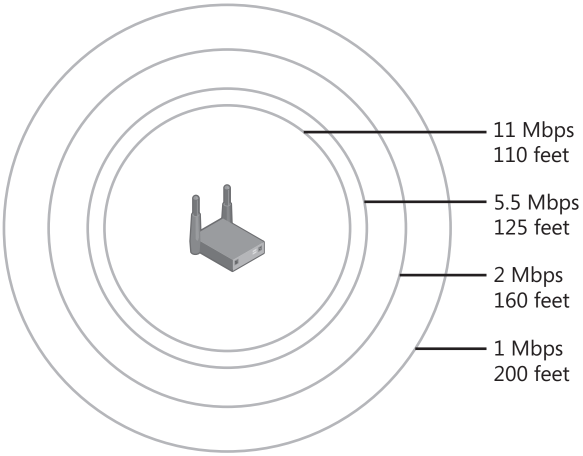

Figure 10-4 Idealized coverage volume example

Signal propagation modifiers change the shape of the ideal coverage volume through radio frequency (RF) attenuation (the reduction of signal strength), shielding, and reflection, which can affect how you deploy your wireless APs. Metal objects within a building or used in the construction of a building can affect the wireless signal. Examples of such objects include:

- Support beams
- Elevator shafts
- Steel reinforcement in concrete
- Heating and air-conditioning ventilation ducts
- Wire mesh that reinforces plaster or stucco in walls
- Walls that contain metal, cinder blocks, and concrete
- Cabinets, metal desks, or other types of large metal equipment

Sources of Interference

Any device that operates on the same frequencies as your wireless devices (in the S-Band ISM, which operates in the frequency range of 2.4 gigahertz [GHz] to 2.5 GHz, or the C-Band ISM, which operates in the frequency range of 5.725 GHz to 5.875 GHz) might interfere with the wireless signals. Sources of interference also change the shape of a wireless AP's ideal coverage volume.

Devices that operate in the S-Band ISM include the following:

- Bluetooth-enabled devices
- Microwave ovens
- 2.4-GHz cordless phones
- Wireless video cameras

- Medical equipment
- Elevator motors

Devices that operate in the C-Band ISM include the following:

- 5-GHz cordless phones
- Wireless video cameras
- Medical equipment

Number of Wireless APs

To determine how many wireless APs to deploy, follow these guidelines:

- Include enough wireless APs to ensure that wireless users have sufficient signal strength from anywhere in the coverage volume.

  Typical wireless APs use antennas that produce a vertically flattened sphere of signal that propagates across the floor of a building. Wireless APs typically have indoor coverage within a 200-foot radius. Include enough wireless APs to ensure signal overlap between the wireless APs.
- Determine the maximum number of simultaneous wireless users per coverage volume.
- Estimate the data throughput that the average wireless user requires. If needed, add more wireless APs, which will:
  - Improve wireless client network bandwidth capacity.
  - Increase the number of wireless users supported within a coverage area.
  - Based on the total data throughput of all users, determine the number of users who can connect to a wireless AP. Obtain a clear picture of throughput before deploying the network or making changes. Some wireless vendors provide an 802.11 simulation tool, which you can use to model traffic in a network and view throughput levels under various conditions.
  - Ensure redundancy in case a wireless AP fails.
- When designing wireless AP placement for performance, use the following best practices:
  - Do not overload your wireless APs with too many connected wireless clients. Although most wireless APs can support hundreds of wireless connections, the practical limit is 20 to 25 connected clients. An average of 2 to 4 users per wireless AP is a good average to maximize the performance while still effectively utilizing the wireless LAN.
  - For higher density situations, lower the signal strength of the wireless APs to reduce the coverage area, thereby allowing more wireless APs to fit in a specific space and more wireless bandwidth to be distributed to more wireless clients.

Authentication Infrastructure

The authentication infrastructure exists to:

- Authenticate the credentials of wireless clients.
- Authorize the wireless connection.
- Inform wireless APs of wireless connection restrictions.
- Record the wireless connection creation and termination for accounting purposes.

The authentication infrastructure for protected wireless connections consists of:

- Wireless APs
- RADIUS servers
- Active Directory domain controllers
- Issuing CAs of a PKI (optional)

If you are using a Windows domain as the user account database for verification of user or computer credentials and for obtaining dial-in properties, use Network Policy Server (NPS) in Windows Server 2008. NPS is a full-featured RADIUS server and proxy that is tightly integrated with Active Directory. See Chapter 9 for additional design and planning considerations for NPS-based RADIUS servers.

NPS performs the authentication of the wireless connection by communicating with a domain controller over a protected remote procedure call (RPC) channel. NPS performs authorization of the connection attempt through the dial-in properties of the user or computer account and network policies configured on the NPS server.

By default, NPS logs all RADIUS accounting information in a local log file (*%SystemRoot%*\System32\Logfiles\***Logfile***.log by default) based on settings configured in the Accounting node in the Network Policy Server snap-in.

Best Practices for Authentication Infrastructure

Best practices to follow for the authentication infrastructure are the following:

- To better manage authorization for wireless connections, create a universal group in Active Directory for wireless access that contains global groups for the user and computer accounts that are allowed to make wireless connections. For example, create a universal group named WirelessAccounts that contains the global groups based on your organization's regions or departments. Each global group contains allowed user and computer accounts for wireless access. When you configure your NPS policies for wireless connections, specify the WirelessAccounts group name.

- From the NPS node of the Network Policy Server snap-in, use the Configure 802.1X Wizard to create a set of policies for 802.1X-authenticated wireless connections. For example, create a set of policies for wireless clients that are members of a specific group and to use a specific authentication method.

Wireless Clients

A Windows-based wireless client is one that is running Windows Server 2008, Windows Vista, Windows XP with Service Pack 2, or Windows Server 2003. You can configure wireless connections on Windows-based wireless clients in the following ways:

- **Group Policy** The Wireless Network (IEEE 802.11) Policies Group Policy extension is part of a Computer Configuration Group Policy Object that can specify wireless network settings in an Active Directory environment.
- **Command line** You can configure wireless settings by using Netsh.exe (running the command **netsh wlan** with the desired parameters). These commands apply only to wireless clients running Windows Vista or Windows Server 2008.

> **Note** To run **netsh wlan** commands on computers running Windows Server 2008, you must add the Wireless LAN Service feature with the Server Manager tool.

- **Wireless XML profiles** Wireless Extensible Markup Language (XML) profiles are XML files that contain wireless network settings. You can use either the Netsh tool or the Wireless Network (IEEE 802.11) Policies Group Policy extension to export and import XML-based wireless profiles.
- **Manually** For a Windows Vista–based or Windows Server 2008–based wireless client, connect to the wireless network when prompted or use the Connect to a Network Wizard from the Network and Sharing Center. For a Windows XP with SP2–based or Windows Server 2003–based wireless client, connect to the wireless network when prompted, or use the Wireless Network Setup Wizard from the Network Connections folder.

Wireless Network (IEEE 802.11) Policies Group Policy Extension

To automate the configuration of wireless network settings for Windows wireless client computers, Windows Server 2008 and Windows Server 2003 Active Directory domains support a Wireless Network (IEEE 802.11) Policies Group Policy extension. This extension allows you to configure wireless network settings as part of Computer Configuration Group Policy for a domain-based Group Policy Object. By using the Wireless Network (IEEE 802.11) Policies Group Policy extension, you can specify a list of preferred networks and their settings to automatically configure wireless LAN settings for wireless clients running Windows Server 2008, Windows Vista, Windows XP with SP2, Windows XP with SP1, or Windows Server 2003.

For each preferred network, you can specify the following:

- Connection settings, such as the wireless network name and whether the wireless network is a non-broadcast network
- Security settings, such as the authentication and encryption method, the EAP type, and the authentication mode
- Advanced 802.1X security settings, such as Single Sign On (for Windows Server 2008 and Windows Vista wireless clients)

These settings are automatically applied to wireless clients running Windows Server 2008, Windows Vista, Windows XP with SP2, and Windows Server 2003 that are members of a Windows Server 2008 or Windows Server 2003 Active Directory domain. You can configure wireless policies by using the Computer Configuration\Windows Settings\Security Settings\Wireless Network (IEEE 802.11) Policies node in the Group Policy Management Editor snap-in.

> **Note** To modify Group Policy settings from a computer running Windows Server 2008, you might need to install the Group Policy Management feature using the Server Manager tool.

By default, there are no Wireless Network (IEEE 802.11) policies. To create a new policy for a Windows Server 2008–based Active Directory domain, right-click Wireless Network (IEEE 802.11) Policies in the Group Policy Management Editor snap-in console tree, and then click Create A New Windows Vista Policy or Create A New Windows XP Policy. For each type of policy, you can create only a single policy. A Windows XP Policy can contain profiles with settings for multiple wireless networks, and each network must have a unique SSID. A Windows Vista policy can also contain profiles with settings for multiple wireless networks with unique SSIDs. Additionally, different profiles can contain multiple instances of the same SSID, each with unique settings. This allows you to configure profiles for mixed-mode deployments in which some clients are using different security technologies, such as WPA and WPA2.

The Windows Vista–based wireless policy contains policy settings specific to Windows Server 2008 and Windows Vista wireless clients. If both types of wireless policies are configured, Windows XP with SP2–based and Windows Server 2003–based wireless clients will use only the Windows XP policy settings, and the Windows Server 2008 and Windows Vista wireless clients will use only the Windows Vista policy settings. If there are no Windows Vista policy settings, Windows Server 2008 and Windows Vista wireless clients will use the Windows XP policy settings.

Windows Vista Wireless Policy The properties dialog box of a Windows Vista wireless policy consists of a General tab and a Network Permissions tab. Figure 10-5 shows the General tab.

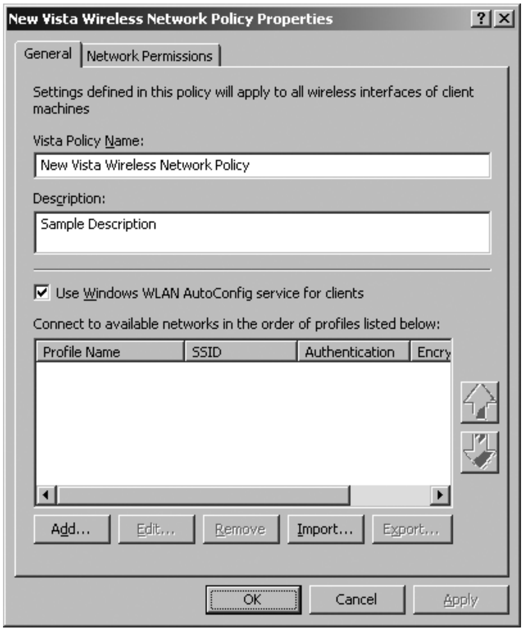

New Vista Wireless Network Policy Properties

General | Network Permissions

Settings defined in this policy will apply to all wireless interfaces of client machines

Vista Policy Name: New Vista Wireless Network Policy

Description: Sample Description

☑ Use Windows WLAN AutoConfig service for clients

Connect to available networks in the order of profiles listed below:

Profile Name	SSID	Authentication	Encry

Figure 10-5 The General tab of a Windows Vista wireless policy

On the General tab, you can configure a name and description for the policy, specify whether to enable the WLAN AutoConfig service (Wireless Auto Configuration), and configure the list of wireless networks and their settings (known as *profiles*) in preferred order. On the General tab, you can import and export profiles as files in XML format. To export a profile to an XML file, select the profile and click Export. To import an XML file as a wireless profile, click Import, and then specify the file's location.

Figure 10-6 shows the Network Permissions tab for a Windows Vista wireless network policy.

The Network Permissions tab is new for Windows Server 2008 and Windows Vista and allows you to specify wireless networks by name that are either allowed or denied access. For example, you can create allow or deny lists.

With an allow list, you can specify the set of wireless networks by name to which a Windows Server 2008 or Windows Vista wireless client is allowed to connect. This is useful for network administrators who want an organization's laptop computers to connect to a specific set of wireless networks, which might include the organization's wireless network in addition to wireless Internet service providers.

With a deny list, you can specify the set of wireless networks by name to which the wireless clients are not allowed to connect. This is useful to prevent managed laptop computers from connecting to other wireless networks that are within range of the organization's wireless network—for example, when an organization occupies a floor of a building and there are other wireless networks of other organization on adjoining floors—or to prevent managed laptop computers from connecting to known unsecured wireless networks.

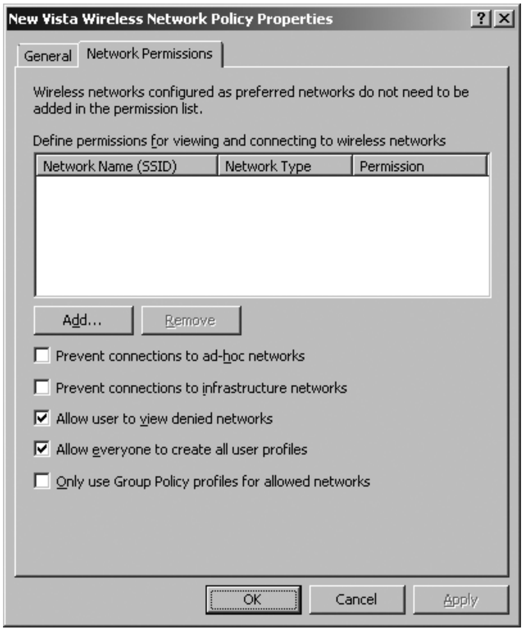

Wireless networks configured as preferred networks do not need to be added in the permission list.

Define permissions for viewing and connecting to wireless networks

Network Name (SSID)	Network Type	Permission

☐ Prevent connections to ad-hoc networks

☐ Prevent connections to infrastructure networks

☑ Allow user to view denied networks

☑ Allow everyone to create all user profiles

☐ Only use Group Policy profiles for allowed networks

Figure 10-6 The Network Permissions tab of a Windows Vista wireless policy

On the Network Permissions tab, there are also settings to prevent connections to either ad-hoc or infrastructure mode wireless networks, to allow the user to view the wireless networks in the list of available networks that have been configured as denied, and to allow any user to create an all-user profile. An *all-user profile* can be used to connect to a specific wireless network by any user with an account on the computer. If this setting is disabled, only users in the Domain Admins or Network Configuration Operators groups can create all-user wireless profiles on the computer. Last, there is a setting to require that the wireless client use Group Policy–based profiles for allowed profiles, rather than local profiles of the same name.

To manage a wireless network profile, in the New Windows Vista Wireless Policy Properties dialog box, on the General tab, either select an existing profile and click Edit, or click Add and then specify whether the new wireless profile is for an infrastructure or ad-hoc mode wireless network. The profile properties dialog box of a Windows Vista wireless network profile consists of a Connection tab and a Security tab. Figure 10-7 shows the default Connection tab for a Windows Vista wireless network profile.

On the Connection tab, you can configure a name for the profile and a list of wireless network names to which this profile applies. You can add new names by typing the name in the Network Name(s) (SSID) box and clicking Add. You can also specify whether the wireless client using this profile will automatically attempt to connect to the wireless networks named in the profile when in range (subject to the preference order of the list of wireless profiles on the General tab for the Windows Vista policy), whether to automatically disconnect from this wireless network if a more preferred wireless network comes within range, and to indicate that the wireless networks in this profile are non-broadcast networks (also known as hidden networks).

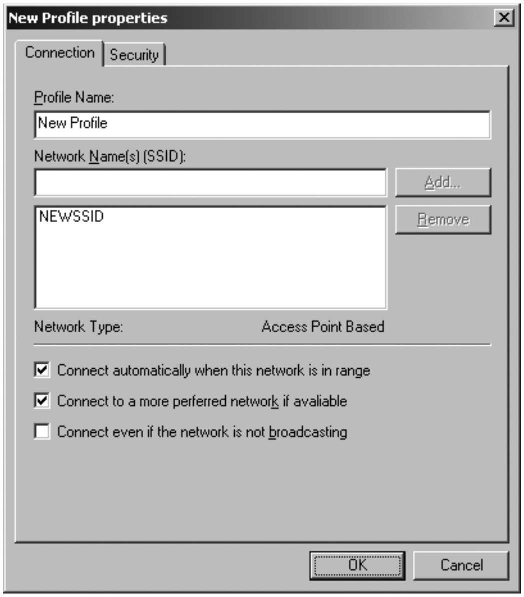

Profile Name: New Profile

Network Name(s) (SSID):

NEWSSID

Network Type: Access Point Based

☑ Connect automatically when this network is in range

☑ Connect to a more perferred network if avaliable

☐ Connect even if the network is not broadcasting

Figure 10-7 The Connection tab for a Windows Vista wireless network profile

Figure 10-8 shows the Security tab for a Windows Vista wireless network profile.

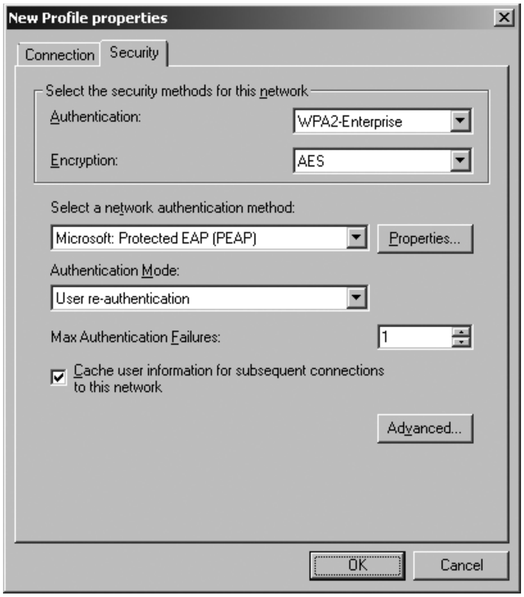

Select the security methods for this network

Authentication: WPA2-Enterprise

Encryption: AES

Select a network authentication method: Microsoft: Protected EAP (PEAP)

Authentication Mode: User re-authentication

Max Authentication Failures: 1

☑ Cache user information for subsequent connections to this network

Figure 10-8 The Security tab for a Windows Vista wireless network profile

On the Security tab, you can configure the authentication and encryption methods for the wireless networks in the profile. For authentication methods, you can select Open, Shared, Wi-Fi Protected Access (WPA)–Personal, WPA-Enterprise, WPA2-Personal, WPA2-Enterprise, and Open with 802.1X. For encryption methods, you can select Wired Equivalent Privacy (WEP), Temporal Key Integrity Protocol (TKIP), and Advanced Encryption Standard (AES). The choice of encryption methods depends on your choice of authentication method.

If you select Open with 802.1X, WPA-Enterprise, or WPA2-Enterprise as the authentication method, you can also configure the network authentication method (the EAP type), the authentication mode (user reauthentication, computer authentication, user authentication, or guest authentication), the number of times authentication attempts can fail before authentication is abandoned, and whether to cache user information for subsequent connections. If you configure this last setting not to cache the user information, when the user logs off, the user credential data is removed from the registry. The result is that when the next user logs on, that user will be prompted for credentials (such as user name and password).

Direct from the Source: Locations of Cached Credentials

For wireless clients running Windows Server 2008 or Windows Vista, the cached credentials are stored at:

HKEY_CURRENT_USER\Software\Microsoft\Wlansvc\UserData\Profiles\Profile-GUID\MSMUserdata

For wireless clients running Windows XP or Windows Server 2003, the cached credentials are stored at:

HKEY_CURRENT_USER\Software\Microsoft\Eapol\UserEapInfo

Clay Seymour, Support Escalation Engineer

Enterprise Platform Support

To configure advanced security settings for the WPA-Enterprise, WPA2-Enterprise, or Open with 802.1X authentication methods, in the New Profile Properties Dialog Box, on the Security tab, click Advanced. Figure 10-9 shows the default Advanced Security Settings dialog box.

In the IEEE 802.1X section, there are settings to specify the number of successive EAP over LAN (EAPOL)-Start messages that are sent out when no response to the initial EAPOL-Start messages is received, the time interval between the retransmission of EAPOL-Start messages when no response to the previously sent EAPOL-Start message is received, the period for which the authenticating client will not perform any 802.1X authentication activity after it has received an authentication failure indication from the authenticator, and the interval for which the authenticating client will wait before retransmitting any 802.1X requests after end-to-end 802.1X authentication has been initiated.

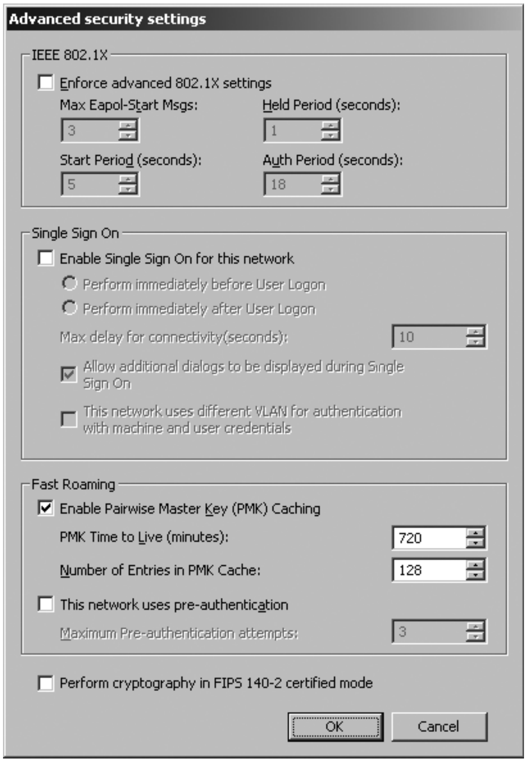

Advanced security settings

IEEE 802.1X

□ Enforce advanced 802.1X settings

Max Eapol-Start Msgs: 3
Held Period (seconds): 1
Start Period (seconds): 5
Auth Period (seconds): 18

Single Sign On

□ Enable Single Sign On for this network

○ Perform immediately before User Logon
○ Perform immediately after User Logon

Max delay for connectivity(seconds): 10

☑ Allow additional dialogs to be displayed during Single Sign On

□ This network uses different VLAN for authentication with machine and user credentials

Fast Roaming

☑ Enable Pairwise Master Key (PMK) Caching

PMK Time to Live (minutes): 720
Number of Entries in PMK Cache: 128

□ This network uses pre-authentication

Maximum Pre-authentication attempts: 3

□ Perform cryptography in FIPS 140-2 certified mode

Figure 10-9 The Advanced Security Settings dialog box

In the Single Sign On section, there are settings to perform wireless authentication immediately before or after the user logon process, specify the number of seconds of delay for connectivity before the user logon process begins, choose whether to prompt the user for additional dialog boxes, and choose whether the wireless networks for this profile use a different virtual LAN (VLAN) for computer or user authentication and to perform a DHCP renewal when switching from the computer-authenticated VLAN to the user-authenticated VLAN. For information about when to use Single Sign On, see "Wireless Authentication Modes" earlier in this chapter.

In the Fast Roaming section, you can configure Pairwise Master Key (PMK) caching and preauthentication options. The Fast Roaming section appears only when you select WPA2-Enterprise as the authentication method on the Security tab. With PMK caching, wireless clients and wireless APs cache the results of 802.1X authentications. Therefore, access is much faster when a wireless client roams back to a wireless AP to which the client already authenticated. You can configure a maximum time to keep an entry in the PMK cache and the maximum number of entries. With preauthentication, a wireless client can perform an 802.1X authentication with other wireless APs in its range while it is still connected to its current wireless AP. If the wireless client roams to a wireless AP with which it has preauthenticated, access time is substantially decreased. You can configure the maximum number of times to attempt preauthentication with a wireless AP.

Note Fast roaming for WPA2 is different than fast reconnect. Fast reconnect minimizes the connection delay in wireless environments when a wireless client roams from one wireless AP to another when using PEAP. With fast reconnect, the Network Policy Server service caches information about the PEAP TLS session so that when reauthenticating, the wireless client does not need to perform PEAP authentication, only MS-CHAP v2 (for PEAP-MS-CHAP v2) or TLS (for PEAP-TLS) authentication. Fast reconnect is enabled by default for Windows wireless clients and for NPS network policies.

A final check box allows you to specify whether to perform AES encryption in a Federal Information Processing Standard (FIPS) 140-2 certified mode. FIPS 140-2 is a U.S. government computer security standard that specifies design and implementation requirements for cryptographic modules. Windows Server 2008 and Windows Vista are FIPS 140-2 certified. When you enable FIPS 140-2 certified mode, Windows Server 2008 or Windows Vista will perform the AES encryption in software, rather than relying on the wireless network adapter. This check box only appears when you select WPA2-Enterprise as the authentication method on the Security tab.

Windows XP Wireless Policy To create a new Windows XP wireless policy, in the Group Policy Management Editor snap-in, in the console tree, right-click Wireless Network (IEEE 802.11) Policies, and then click Create A New Windows XP Policy. The properties dialog box of a Windows XP wireless policy consists of a General tab and a Preferred Networks tab.

Figure 10-10 shows the General tab for a Windows XP wireless network policy.

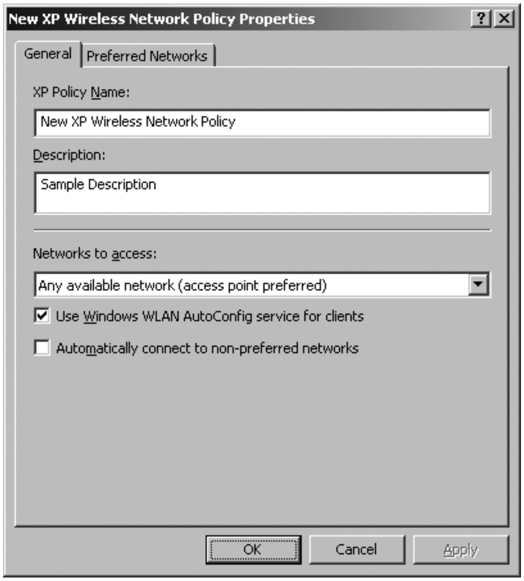

XP Policy Name: New XP Wireless Network Policy

Description: Sample Description

Networks to access: Any available network (access point preferred)

☑ Use Windows WLAN AutoConfig service for clients

☐ Automatically connect to non-preferred networks

Figure 10-10 The General tab for a Windows XP wireless network policy

On the General tab, you can configure a name and description for the policy, specify whether the Wireless Zero Configuration service is enabled, select the types of wireless networks to access (any available, infrastructure, or ad-hoc networks), and specify whether to automatically connect to non-preferred networks.

Figure 10-11 shows the Preferred Networks tab for a Windows XP wireless policy.

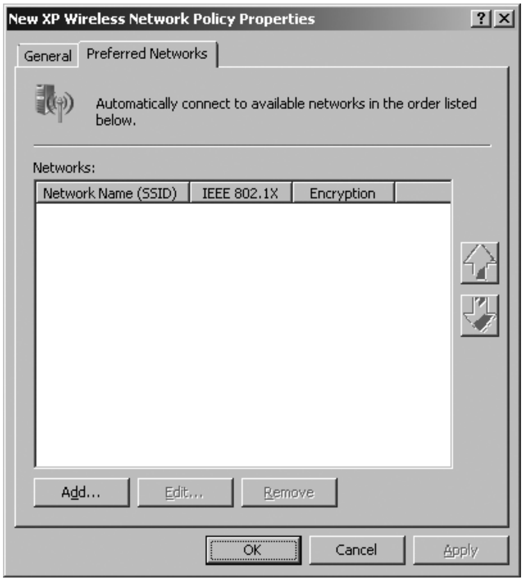
New XP Wireless Network Policy Properties

Automatically connect to available networks in the order listed below.

Networks:

Network Name (SSID)	IEEE 802.1X	Encryption

Figure 10-11 The Preferred Networks tab for a Windows XP wireless policy

On the Preferred Networks tab, you can manage the list of preferred wireless networks and their order of preference. To manage a wireless network profile from the Preferred Networks tab of the Windows XP wireless policy properties dialog box, either select an existing profile and click Edit, or click Add and then specify whether the new wireless profile is for an infrastructure or ad-hoc mode wireless network. The properties dialog box of a preferred wireless network consists of a Network Properties tab and an IEEE 802.1X tab.

Figure 10-12 shows the Network Properties tab for a preferred wireless infrastructure network.

On the Network Properties tab, you can add a description for the preferred network, specify whether the wireless network is a non-broadcast network (infrastructure), select the authentication and encryption methods, and, for WPA2, configure advanced fast roaming settings.

Figure 10-13 shows the default IEEE 802.1X tab for a preferred wireless network.

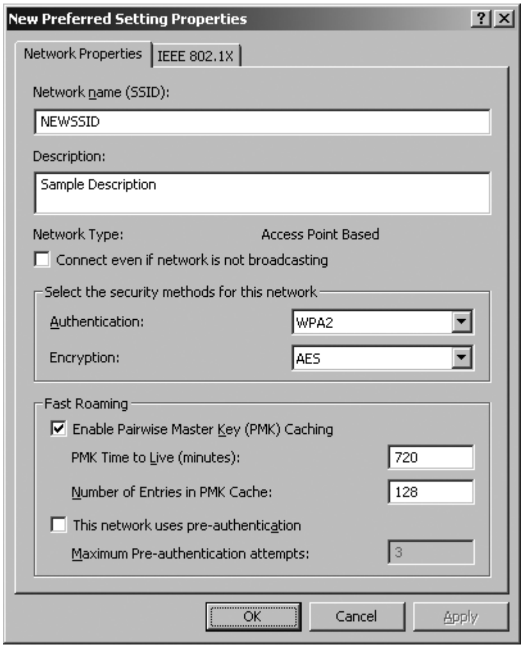

Network name (SSID): NEWSSID

Description: Sample Description

Network Type: Access Point Based

□ Connect even if network is not broadcasting

Select the security methods for this network

Authentication: WPA2

Encryption: AES

Fast Roaming

☑ Enable Pairwise Master Key (PMK) Caching

PMK Time to Live (minutes): 720

Number of Entries in PMK Cache: 128

□ This network uses pre-authentication

Maximum Pre-authentication attempts: 3

Figure 10-12 The Network Properties tab for a preferred wireless infrastructure network

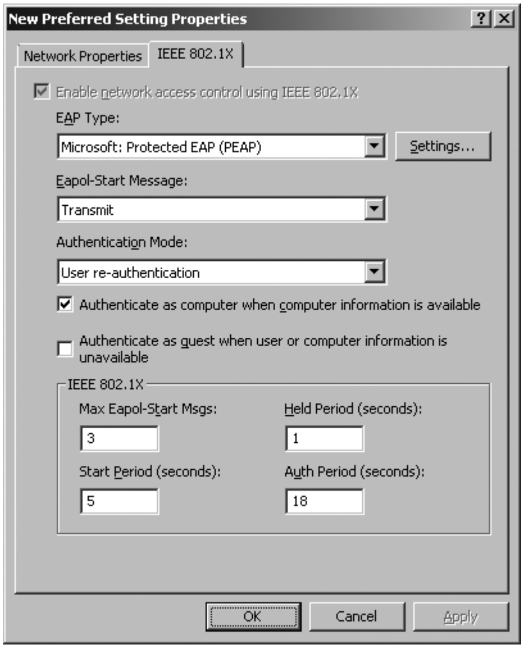

☑ Enable network access control using IEEE 802.1X

EAP Type: Microsoft: Protected EAP (PEAP)

Eapol-Start Message: Transmit

Authentication Mode: User re-authentication

☑ Authenticate as computer when computer information is available

□ Authenticate as guest when user or computer information is unavailable

IEEE 802.1X

Max Eapol-Start Msgs: 3

Held Period (seconds): 1

Start Period (seconds): 5

Auth Period (seconds): 18

Figure 10-13 The IEEE 802.1X tab for a preferred wireless network

On the IEEE 802.1X tab, you can specify the EAP type and configure its settings, specify when to send the EAPOL-Start message, choose the authentication mode, specify whether to authenticate with computer credentials or as a guest, and set advanced 802.1X settings.

Command-Line Configuration

Windows Vista supports a command-line interface that allows you to configure some of the wireless settings that are available from the wireless dialog boxes in the Network Connections folder or through the Wireless Network (IEEE 802.11) Policies Group Policy extension. Command-line configuration of wireless settings can help deployment of wireless networks in the following situations:

- **Automated script support for wireless settings without using Group Policy** The Wireless Network (IEEE 802.11) Policies Group Policy extension applies only in an Active Directory domain. For an environment without a Group Policy infrastructure, a script that automates the configuration of wireless connections can be run either manually or automatically, such as part of the logon script.
- **Bootstrap of a wireless client onto the organization's protected wireless network** A wireless client computer that is not a member of the domain cannot connect to the organization's protected wireless network. Additionally, a computer cannot join the domain until it has successfully connected to the organization's protected wireless network. A command-line script provides a method to connect to the organization's secure wireless network to join the domain.

To perform command-line configuration of Windows Vista–based and Windows Server 2008–based wireless clients, run the **netsh wlan** command with the appropriate parameters.

> **More Info** For more information about **netsh wlan** command syntax, see Netsh Commands for Wireless Local Area Network (WLAN) at *http://go.microsoft.com/fwlink/?LinkID=81751*.

XML-Based Wireless Profiles

To simplify command-line configuration of Windows Vista or Windows Server 2008 wireless clients, you can export the configuration of a wireless profile to an XML file that can be imported on other wireless clients. You can export a wireless profile from a wireless client by running the **netsh wlan export profile** command or by using the General tab of the Windows Vista wireless policy properties dialog box. To import a wireless profile, run **netsh wlan add profile**.

Design Choices for Wireless Clients

The design choices for wireless clients are the following:

- To prevent your Windows Vista or Windows Server 2008 wireless clients from connecting to certain wireless networks, configure a list of denied wireless networks on the Network Permissions tab of the Windows Vista wireless policy properties dialog box, or run the **netsh wlan add filter** command.
- To configure your Windows Vista or Windows Server 2008 wireless clients to connect to only specific wireless networks, configure a list of allowed wireless networks on the Network Permissions tab of the Windows Vista wireless policy dialog box, or run the **netsh wlan add filter** command.

Requirements for Wireless Clients

The requirements for wireless clients are the following:

- To use WPA2, wireless clients must be running Windows XP with SP2 and the Wireless Client Update for Windows XP with Service Pack 2, Windows Vista, or Windows Server 2008.
- Command-line configuration using the **netsh wlan** command, export and import of wireless XML profiles, and Single Sign On are supported by wireless clients running only Windows Vista or Windows Server 2008.
- To deploy 802.1X enforcement with Network Access Protection, you must configure your wireless clients to use a PEAP-based authentication method.

Best Practices for Wireless Clients

Best practices for wireless clients are the following:

- For a small number of wireless clients, configure each wireless client manually.
- For enterprise deployment of wireless configuration in an Active Directory environment, use the Wireless Network (IEEE 802.11) Wireless Policies Group Policy extension.
- For enterprise deployment of wireless configuration through the use of scripts, create wireless XML profiles and configure wireless clients with a script containing the **netsh wlan add profile** command.

PKI

To perform authentication for wireless connections using PEAP-TLS or EAP-TLS, a PKI must be in place to issue computer or user certificates to wireless clients and computer certificates to RADIUS servers. For PEAP-MS-CHAP v2–based authentication, a PKI is not required. It is possible to purchase certificates from a third-party CA to install on your NPS servers. You might also need to distribute the root CA certificate of third-party computer certificates to your wireless client computers.

PKI for Smart Cards

The use of smart cards for user authentication is the strongest form of user authentication in Windows. For wireless connections, you can use smart cards with the EAP-TLS or PEAP-TLS authentication method. The individual smart cards are distributed to users who have a computer with a smart card reader. To log on to the computer, you must insert the smart card into the smart card reader and type the smart card personal identification number (PIN). When the user attempts to make a wireless connection, the smart card certificate is sent during the connection negotiation process.

PKI for User Certificates

User certificates that are stored in the Windows registry for user authentication can be used in place of smart cards. However, it is not as strong a form of authentication. With smart cards, the user certificate issued during the authentication process is made available only when the user possesses the smart card and has knowledge of the PIN to log on to the computer. With user certificates, the user certificate issued during the authentication process is made available when the user logs on to the computer using a domain-based user name and password. Just as with smart cards, authentication using user certificates for wireless connections uses the EAP-TLS or PEAP-TLS authentication methods.

To deploy user certificates in your organization, first deploy a PKI. You'll then need to install a user certificate for each user. The easiest way to accomplish this is if Windows Certificate Services is installed as an enterprise CA. Then configure Group Policy settings for user certificate autoenrollment. For more information, see the section titled "Deploying Certificates" later in this chapter.

When the wireless client attempts user-level authentication for a wireless connection, the wireless client computer sends the user certificate during the authentication process.

PKI for Computer Certificates

Computer certificates are stored in the Windows registry for computer-level authentication for wireless access with the EAP-TLS or PEAP-TLS authentication methods. To deploy computer certificates in your organization, first deploy a PKI. You'll then need to install a computer certificate for each computer. The easiest way to accomplish this is if Windows Active Directory Certificate Services or Certificate Services is installed as an enterprise CA. Then, configure Group Policy settings for computer certificate autoenrollment. For more information, see "Deploying Certificates" later in this chapter.

When the wireless client attempts computer-level authentication for a wireless connection, the wireless client computer sends the computer certificate during the authentication process.

Requirements for PKI

Requirements for PKI for a protected wireless network are the following:

- For computer-level authentication with EAP-TLS or PEAP-TLS, you must install computer certificates, also known as *machine* certificates, on each wireless client.

  The computer certificates of the wireless clients must be valid and verifiable by the NPS servers; the NPS servers must have a root CA certificate for the CA that issued the computer certificates of the wireless client.

- For user-level authentication with EAP-TLS or PEAP-TLS, you must use a smart card, or you must install a user certificate on each wireless client.

  The smart card or user certificates of the wireless clients must be valid and verifiable by the NPS servers; the NPS servers must have the root CA certificates of the issuing CAs of the smart card or user certificates of the wireless clients.

- You must install the root CA certificates of the issuing CA of the NPS server computer certificates on each wireless client.

  The computer certificates of the NPS servers must be valid and verifiable by each wireless client; the wireless clients must have a root CA certificates for the CAs that issued the computer certificates of the NPS servers.

- For EAP-TLS authentication, the requirements for the user certificate, smart card certificate, or computer certificate of the wireless client are as follows:
  - The certificate must contain a private key.
  - The certificate must be issued by an enterprise CA or mapped to a user or computer account in Active Directory.
  - The certificate must be chained to a trusted root CA on the NPS server and must not fail any of the checks that are performed by CryptoAPI and specified in the network policy for wireless connections.
  - The certificate must be configured with the Client Authentication purpose in the Enhanced Key Usage field (the object identifier for Client Authentication is 1.3.6.1.5.5.7.3.2).
  - The Subject Alternative Name field must contain the user principal name (UPN) of the user or computer account.
- For EAP-TLS authentication, the requirements for the computer certificate of the NPS server are as follows:
  - The certificate must contain a private key.

- ❑ The Subject field must contain a value.
- ❑ The certificate must be chained to a trusted root CA on the wireless clients and must not fail any of the checks that are performed by CryptoAPI and specified in the network policy for wireless connections.
- ❑ The certificate must be configured with the Server Authentication purpose in the Enhanced Key Usage field (the object identifier for Server Authentication is 1.3.6.1.5.5.7.3.1).
- ❑ The certificate must be configured with a required cryptographic service provider (CSP) value of Microsoft RSA SChannel Cryptographic provider.
- ❑ The Subject Alternative Name field of the certificate, if used, must contain the DNS name of the NPS server.

Best Practices for PKI

Best practices for the PKI for protected wireless access are the following:

- For computer certificates with EAP-TLS or PEAP-TLS, if you are using a Windows Server 2008 enterprise CA as an issuing CA, configure your Active Directory domain for autoenrollment of computer certificates using a Computer Configuration Group Policy. Each computer that is a member of the domain automatically requests a computer certificate when the Computer Configuration Group Policy is updated.
- For registry-based user certificates for EAP-TLS or PEAP-TLS, if you are using a Windows Server 2008 enterprise CA as an issuing CA, use a User Configuration Group Policy to configure your Active Directory domain for autoenrollment of user certificates. Each user who successfully logs on to the domain automatically requests a user certificate when the User Configuration Group Policy is updated.
- If you have purchased third-party computer certificates for your NPS servers for PEAP-MS-CHAP v2 authentication, and the wireless clients do not have the root CA certificate of the issuing CA of the NPS server computer certificates installed, use Group Policy to install the root CA certificate of the issuing CA of the NPS server computer certificates on your wireless clients. Each computer that is a member of the domain automatically receives and installs the root CA certificate when the Computer Configuration Group Policy is updated.
- For EAP-TLS, PEAP-TLS, and PEAP-MS-CHAP v2 authentication, it is possible to configure the wireless clients so that they do not validate the certificate of the NPS server. If so, it is not required to have computer certificates on the NPS servers and their root CA certificates on wireless clients. However, having the wireless clients validate the certificate of the NPS server is recommended for mutual authentication of the wireless client and NPS server. With mutual authentication, you can protect your wireless clients from connecting to rogue wireless APs with spoofed authentication servers.

802.1X Enforcement with NAP

NAP for Windows Server 2008, Windows Vista, and Windows XP with Service Pack 3 provides components and an application programming interface (API) set that help you enforce compliance with health policies for network access or communication. Developers and network administrators can create solutions for validating computers that connect to their networks, can provide needed updates or access to needed resources, and can limit the access of noncompliant computers.

802.1X Enforcement is one of the NAP enforcement methods included with Windows Server 2008, Windows Vista, and Windows XP. With 802.1X Enforcement, an 802.1X-authenticated wireless client must prove that it is compliant with system health requirements before being allowed full access to the intranet. If the wireless client is not compliant with system health requirements, the wireless AP places the wireless client on a restricted network containing servers that have resources to bring the wireless client back into compliance. The wireless AP enforces the restricted access through packet filters or a VLAN ID that are assigned to the wireless connection. After correcting its health state, the wireless client validates its health state again, and if compliant, the constraints on the wireless connection that confine the access to the restricted network are removed.

In order for 802.1X Enforcement to work, you must already have a working protected wireless deployment that uses a PEAP-based authentication method. For the details on deploying 802.1X Enforcement after successfully deploying a protected wireless network solution, see Chapter 17.

Deploying Protected Wireless Access

To deploy a protected wireless network using Windows Server 2008 and Windows Vista, follow these steps:

1. Deploy certificates.
2. Configure Active Directory for user accounts and groups.
3. Configure NPS servers.
4. Deploy wireless APs.
5. Configure wireless clients.

Deploying Certificates

Each wireless client in the following authentication configurations needs a computer certificate:

- **Computer authentication with EAP-TLS or PEAP-TLS and computer certificates** Each wireless client computer needs a computer certificate.

- **User authentication with EAP-TLS or PEAP-TLS and either smart cards or registry-based user certificates** Each wireless user needs a smart card, or each wireless client computer needs a user certificate.
- **User or computer authentication with PEAP-MS-CHAP v2** Each wireless client needs the root CA of the issuing CA of the NPS server's computer certificate.

Deploying Computer Certificates

To install computer certificates for EAP-TLS or PEAP-TLS authentication, a PKI must be present to issue certificates. Once the PKI is in place, you can install a computer certificate on wireless clients and NPS servers in the following ways:

- By configuring autoenrollment of computer certificates to computers in an Active Directory domain (recommended)
- By using the Certificates snap-in to request a computer certificate
- By using the Certificates snap-in to import a computer certificate
- By executing a CAPICOM script that requests a computer certificate

For more information, see "Deploying PKI" in Chapter 9.

Deploying User Certificates

You can install a user certificate on wireless clients in the following ways:

- By configuring autoenrollment of user certificates to users in an Active Directory domain (recommended)
- By using the Certificates snap-in to request a user certificate
- By using the Certificates snap-in to import a user certificate
- By requesting a certificate over the Web
- By executing a CAPICOM script that requests a user certificate

For more information, see "Deploying PKI" in Chapter 9.

Deploying Root CA Certificates

If you use PEAP-MS-CHAP v2 authentication, you might need to install the root CA certificates of the computer certificates that are installed on your NPS servers on your wireless clients. If the root CA certificate of the issuer of the computer certificates that are installed on the NPS servers is already installed as a root CA certificate on your wireless clients, no other configuration is necessary. For example, if your root CA is a Windows Server 2008–based online root enterprise CA, the root CA certificate is automatically installed on each domain member computer through Group Policy.

To verify whether the correct root CA certificate is installed on your wireless clients, you need to determine:

- The root CA of the computer certificates installed on the NPS servers
- Whether a certificate for the root CA is installed on your wireless clients

To Determine the Root CA of the Computer Certificates Installed on the NPS Servers

1. In the console tree of the Certificates snap-in for the NPS server computer account, expand Certificates (Local Computer or *Computername*), expand Personal, and then click Certificates.
2. In the details pane, double-click the computer certificate that is being used by the NPS server for PEAP-MS-CHAP v2 authentication.
3. In the Certificate properties dialog box, on the Certification Path tab, note the name at the top of the certification path. This is the name of the root CA.

To Determine Whether a Certificate for the Root CA Is Installed on Your Wireless Client

1. In the console tree of the Certificates snap-in for the wireless client computer account, expand Certificates (Local Computer or *Computername*), expand Trusted Root Certification Authorities, and then click Certificates.
2. Examine the list of certificates in the details pane for a name matching the root CA for the computer certificates issued to the NPS servers.

You must install the root CA certificates of the issuers of the computer certificates of the NPS servers on each wireless client that does not contain them. The easiest way to install a root CA certificate on all your wireless clients is through Group Policy. For more information, see “Deploying PKI” in Chapter 9.

Configuring Active Directory for Accounts and Groups

To configure Active Directory for wireless access, do the following for the user and computer accounts that will be used to authenticate wireless connections:

- On the Dial-in tab, set the network access permission to Allow Access or Control Access Through NPS Network Policy. With this setting, the permission for access to the network is set by the Access Permission in the NPS network policy. By default, in native-mode domains, new user accounts and computer accounts have the network access permission set to Control Access Through NPS Network Policy.
- Organize the computer and user accounts into the appropriate universal and global groups to take advantage of group-based network policies.

Configuring NPS Servers

Configure and deploy your NPS servers as described in Chapter 9, taking the following steps:

1. Install a computer certificate on each NPS server.
2. Install the root CA certificates of the computer or user certificates of the wireless clients on each NPS server (if needed).
3. Configure logging on the primary NPS server.
4. Add RADIUS clients to the primary NPS server corresponding to each wireless AP.
5. Create on the primary NPS server a set of policies that are customized for wireless connections using the universal group name for your wireless accounts.

For the details of steps 1–4, see Chapter 9.

To Create a Set of Policies for Wireless Connections

1. In the console tree of the Network Policy Server snap-in, click NPS.
2. In the details pane, under Standard Configuration, select RADIUS Server For 802.1X Wireless Or Wired Connections from the drop-down list, and then click Configure 802.1X.
3. In the Configure 802.1X Wizard, on the Select 802.1X Connections Type page, click Secure Wireless Connections from the drop-down list, and then in the Policy Name box, type a name (or use the name created by the wizard). Click Next.
4. On the Specify 802.1X Switches page, add RADIUS clients as needed that correspond to your wireless APs. Click Next.
5. On the Configure An Authentication Method page, configure the EAP type to use for wireless connections.

   To configure EAP-TLS, in the Type drop-down list, select Microsoft: Smart Card Or Other Certificate, and then click Configure. In the Smart Card Or Other Certificate Properties dialog box, select the computer certificate to use for wireless connections, and then click OK. If you cannot select the certificate, the cryptographic service provider for the certificate does not support Secure Channel (SChannel). SChannel support is required for NPS to use the certificate for EAP-TLS authentication.

   To configure PEAP-MS-CHAP v2, in the Type drop-down list, select Protected EAP (PEAP), and then click Configure. In the Edit Protected EAP Properties dialog box, select the computer certificate to use for wireless connections, and then click OK. If you cannot select the certificate, the cryptographic service provider for the certificate does not support SChannel. SChannel support is required for NPS to use the certificate for PEAP authentication.

To configure PEAP-TLS, in the Type drop-down list, select Protected EAP (PEAP), and then click Configure. In the Edit Protected EAP Properties dialog box, select the computer certificate to use for wireless connections. If you cannot select the certificate, the cryptographic service provider for the certificate does not support SChannel. Under EAP Types, click Secured Password (EAP-MSCHAP v2), and then click Remove. Click Add. In the Add EAP dialog box, click Smart Card Or Other Certificate, and then click OK. In the Edit Protected EAP Properties dialog box, under EAP Types, click Smart Card Or Other Certificate, and then click Edit. In the Smart Card Or Other Certificate Properties dialog box, select the computer certificate to use for wireless connections, and then click OK. If you cannot select the certificate, the cryptographic service provider for the certificate does not support Secure Channel (SChannel). Click OK twice.

6. Click Next. On the Specify User Groups page, add the groups containing the wireless computer and user accounts (for example, WirelessAccounts).
7. On the Configure A Virtual LAN (VLAN) page, click Configure if needed to specify the RADIUS attributes and their values that configure your wireless APs for the appropriate VLAN. Click Next.
8. On the Completing New IEEE 802.1X Secure Wired And Wireless Connections And RADIUS Clients page, click Finish.

After you have configured the primary NPS server with the appropriate logging, RADIUS client, and policy settings, copy the configuration to the secondary or other NPS servers. For more information, see Chapter 9.

Deploying Wireless APs

To deploy your wireless APs, do the following:

1. Perform an analysis of wireless AP locations based on plans of floors and buildings.
2. Temporarily install your wireless APs.
3. Perform a site survey analyzing signal strength in all areas.
4. Relocate wireless APs or sources of RF attenuation or interference.
5. Verify the coverage volume.
6. Update the architectural drawings to reflect the final number and placement of the wireless APs.
7. Configure TCP/IP, security, and RADIUS settings.

These steps are discussed in more detail in the following sections.

Note An alternate method of performing a site survey is to move a single wireless AP around to various locations within your site to discover interference issues and identify the eventual locations of your wireless APs. This method allows you to determine the feasibility of a wireless network within your site before you install numerous wireless APs.

Perform an Analysis of Wireless AP Locations

Obtain or create scaled architectural drawings of each floor for each building for which wireless access is being planned. On the drawing for each floor, identify the offices, conference rooms, lobbies, or other areas where you want to provide wireless coverage.

It might be useful to enable wireless coverage for a building in its entirety rather than for specific locations within the building. This type of coverage can prevent connectivity problems that might result from undocking a laptop from an office for use in a different part of your building.

On the plans, indicate the devices that interfere with the wireless signals, and mark the building construction materials or objects that might attenuate, reflect, or shield wireless signals. Then indicate the locations of wireless APs so that each wireless AP is no farther than 200 feet from an adjacent wireless AP.

After you have determined the initial locations of the wireless APs, you must determine their channels and then assign those channel numbers to each wireless AP.

To Select the Channels for the Wireless APs

1. Identify the wireless networks owned by other organizations in the same building. Find out the placement of their wireless APs and the assigned channel.

   Wireless network signal waves travel through floors and ceilings, so wireless APs located near each other on different floors need to be set to non-overlapping channels. If another organization located on a floor adjacent to your organization’s offices has a wireless network, the wireless APs for that organization might interfere with the wireless APs in your network. Contact the other organization to determine the placement and channel numbers of their wireless APs to ensure that your own wireless APs that provide overlapping coverage use a different channel number.

2. Identify overlapping wireless signals on adjacent floors within your own organization.

3. After identifying overlapping coverage volumes outside and within your organization, assign channel numbers to your wireless APs.

To Assign the Channel Numbers to the Wireless APs

1. Assign channel 1 to the first wireless AP.

2. Assign channels 6 and 11 to the wireless APs that overlap coverage volumes with the first wireless AP ensuring that those wireless APs do not also interfere with other coverage volumes with the same channel.

3. Continue assigning channel numbers to the wireless APs ensuring that any two wireless APs with overlapping coverage are assigned different channel numbers that are separated by at least five channels.

Temporarily Install Your Wireless APs

Based on the locations and channel configurations indicated in your plan-based analysis of wireless AP locations, temporarily install your wireless APs.

Perform a Site Survey

Perform a site survey by walking around the building and its floors with a laptop computer equipped with an 802.11 wireless adapter and site survey software. (Site survey software ships with most wireless adapters and wireless APs.) Determine the signal strength and bit rate for the coverage volume for each installed wireless AP.

Relocate Wireless APs or Sources of RF Attenuation or Interference

In locations where signal strength is low, you can make any of the following adjustments to improve the signal:

- Reposition the temporarily installed wireless APs to increase the signal strength for that coverage volume.
- Reposition or eliminate devices that interfere with signal strength (such as Bluetooth devices or microwave ovens).
- Reposition or eliminate metal obstructions that interfere with signal propagation (such as filing cabinets and appliances).
- Add more wireless APs to compensate for the weak signal strength.

> **Note** If you add a wireless AP, you might have to change the channel numbers of adjacent wireless APs.

- Purchase antennas to meet the requirements of your building infrastructure.

For example, to eliminate interference between wireless APs located on adjoining floors in your building, you can purchase directional antennas that flatten the signal (forming a donut-shaped coverage volume) to increase the horizontal range and further decrease the vertical range.

Verify Coverage Volume

Perform another site survey to verify that the changes made to the configuration or placement of the wireless APs eliminated the locations with low signal strength.

Update Your Plans

Update the architectural drawings to reflect the final number and placement of the wireless APs. Indicate the boundaries of the coverage volume and where the data rate changes for each wireless AP.

Configure TCP/IP, Security, and RADIUS Settings

Configure your wireless APs with the following:

- A new wireless network name and strong administrator password
- A static IPv4 address, subnet mask, and default gateway for the wireless subnet on which it is placed
- WPA2 or WPA with 802.1X authentication (WPA2-Enterprise or WPA-Enterprise).

  Configure the following RADIUS settings:

  - The IP address or name of a primary RADIUS server, the RADIUS shared secret, UDP ports for authentication and accounting, and failure detection settings
  - The IP address or name of a secondary RADIUS server, the RADIUS shared secret, UDP ports for authentication and accounting, and failure detection settings

To balance the load of RADIUS traffic between the two NPS servers, configure half of the wireless APs with the primary NPS server as the primary RADIUS server and the secondary NPS server as the secondary RADIUS server. Then, configure the other half of the wireless APs with the secondary NPS server as the primary RADIUS server and the primary NPS server as the secondary RADIUS server.

If the wireless APs require vendor-specific attributes (VSAs) or additional RADIUS attributes, you must add the VSAs or attributes to the wireless network policy of the NPS servers. If you add VSAs or RADIUS attributes to the wireless network policy on the primary NPS server, copy the primary NPS server configuration to the secondary NPS server.

Configuring Wireless Clients

You can configure wireless clients in the following three ways:

- Through Group Policy
- By configuring and deploying wireless XML profiles
- Manually

Configuring Wireless Clients Through Group Policy

To configure Wireless Network (IEEE 802.11) Policies Group Policy settings, perform the following steps:

1. From a computer running Windows Server 2008 that is a member of your Active Directory domain, open the Group Policy Management snap-in.
2. In the console tree, expand Forest, expand Domains, and then click the name of the domain to which your wireless clients belong.

3. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
4. In the console tree of the Group Policy Management Editor snap-in, expand the Group Policy Object, then Computer Configuration, then Windows Settings, then Security Settings, and then Wireless Network (IEEE 802.11) Policies.
5. Right-click Wireless Network (IEEE 802.11) Policies, and then click either Create a New Windows Vista Policy or Create a New Windows XP Policy.

For a new Windows Vista wireless policy, perform the following steps:

1. For the newly created Windows Vista wireless network policy, on the General tab, type a name for the policy and a description.
2. On the Network Permissions tab, add allowed and denied wireless networks by name as needed.
3. On the General tab, click Add to add a wireless network profile, and then click Infrastructure to specify an infrastructure mode wireless network.
4. On the Connection tab, type the wireless network name (SSID) and a description (optional), and then specify connection settings as needed.
5. On the Security tab, specify the authentication and encryption security methods.
   - ❑ For WPA2, in the Authentication section, select WPA2, and then in the Encryption area, select AES.
   - ❑ For WPA, select WPA in Authentication and either TKIP or AES in Encryption. Select AES only if both your wireless clients and wireless APs support WPA with AES encryption.
6. In the Select A Network Authentication Method drop-down list, specify the EAP type.
   - ❑ For EAP-TLS:
     a. Select Smart Card Or Other Certificate, and then click Properties.
     b. In the Smart Card Or Other Certificate Properties dialog box, configure EAP-TLS settings as needed, and then click OK. By default, EAP-TLS uses a registry-based certificate and validates the server certificate.
   - ❑ For PEAP-MS-CHAP v2, no additional configuration is required. PEAP-MS-CHAP v2 is the default authentication method.

   Specify the authentication mode and other settings as needed.
7. To configure advanced settings for 802.1X, including Single Sign On and Fast Roaming, click Advanced and specify settings as needed. Click OK when complete.
8. Click OK twice to save the changes.

For a new Windows XP wireless policy, perform the following steps:

1. For the newly created Windows XP wireless network policy, on the General tab, change settings as needed.
2. On the Preferred Networks tab, click Add to add a preferred network, and then click Infrastructure to specify an infrastructure mode wireless network.
3. On the Network Properties tab, type the wireless network name (SSID), a description (optional), specify whether this wireless network is non-broadcast, and then specify the security methods.
   - For WPA2, in the Authentication drop-down list, select WPA2, and then in the Encryption drop-down list, select AES.
   - For WPA, in the Authentication drop-down list, select WPA, and then in the Encryption drop-down list, select TKIP. Select AES only if both your wireless clients and wireless APs support WPA with AES encryption.
4. On the IEEE 802.1X tab, specify the EAP type.
   - For EAP-TLS:
     a. In the EAP Type drop-down list, select Smart Card Or Other Certificate, and then click Settings.
     b. In the Smart Card Or Other Certificate Properties dialog box, configure EAP-TLS settings as needed, and then click OK. By default, EAP-TLS uses a registry-based certificate and validates the server certificate.
   - For PEAP-MS-CHAP v2, no additional configuration is required. PEAP-MS-CHAP v2 is the default authentication method.
5. Also on the IEEE 802.1X tab, specify the authentication mode and other settings as needed.
6. Click OK twice to save changes.

> **Note** To obtain help information for the dialog boxes of the Wireless Network (IEEE 802.11) Policies Group Policy extension, press the F1 key.

The next time your Windows Server 2008, Windows Vista, Windows XP with SP2, Windows XP with SP1, or Windows Server 2003 wireless clients update the Computer Configuration Group Policy, the wireless network settings in the Group Policy Object will be automatically applied.

Configuring and Deploying Wireless Profiles

You can also manually configure wireless clients running Windows Vista or Windows Server 2008 on a wireless network by importing a wireless profile in XML format by running the **netsh wlan add profile** command. To create an XML-based wireless profile, configure a Windows Vista or Windows Server 2008 wireless client with a wireless network that has all the appropriate settings including the authentication method, encryption methods, and EAP type. Then, run the **netsh wlan export profile** command to write the wireless network profile to an XML file. You can also create, configure, and export an XML profile from a Windows Vista wireless policy.

Manually Configuring Wireless Clients

If you have a small number of wireless clients, you can manually configure wireless connections for each wireless client. For Windows Server 2008 and Windows Vista wireless clients, run the Set Up a Connection Wizard or the Network Wizard. For Windows XP with SP2 wireless clients, run the New Connection Wizard. The following sections describe how to manually configure the EAP-TLS, PEAP-TLS, and PEAP-MS-CHAP v2 authentication methods for Windows wireless clients.

EAP-TLS To manually configure EAP-TLS authentication on a wireless client running Windows Server 2008 or Windows Vista, do the following:

1. In the Network and Sharing Center, click the Manage Wireless Networks task. In the Manage Wireless Networks window, double-click your wireless network name.
2. On the Security tab, in the Security Type box, select WPA-Enterprise or WPA2-Enterprise. In the Choose A Network Authentication Method drop-down list, select Smart Card Or Other Certificate, and then click Settings.
3. In the Smart Card Or Other Certificate Properties dialog box, to use a registry-based user certificate, select Use A Certificate On This Computer. For a smart card–based user certificate, select Use My Smart Card.

   If you want to validate the computer certificate of the NPS server, select Validate Server Certificate (recommended and enabled by default). If you want to specify the names of the NPS servers that must perform the TLS authentication, select Connect To These Servers, and then type the names. Click OK twice.

To manually configure EAP-TLS authentication on a wireless client running Windows XP with SP2, Windows XP with SP1, or Windows Server 2003, do the following:

1. Obtain properties of the wireless connection in the Network Connections folder. On the Wireless Networks tab, in the list of preferred networks, click the name of the wireless network, and then click Properties.
2. On the Authentication tab, select Enable Network Access Control Using IEEE 802.1X and the Smart Card Or Other Certificate EAP type. This is enabled by default.

3. Click Properties. In the properties dialog box of the Smart Card or other Certificate EAP type, to use a registry-based user certificate, select Use A Certificate On This Computer. For a smart card–based user certificate, select Use My Smart Card.

   If you want to validate the computer certificate of the NPS server, select Validate Server Certificate (recommended and enabled by default). If you want to specify the names of the authentication servers that must perform the TLS authentication, select Connect To These Servers, and then type the names.

4. Click OK to save changes to the Smart Card or other Certificate EAP type.

PEAP-TLS To manually configure PEAP-TLS authentication on a wireless client running Windows Server 2008 or Windows Vista, do the following:

1. In the Network and Sharing Center, click the Manage Wireless Networks task. In the Manage Wireless Networks window, double-click your wireless network name.
2. On the Security tab, in the Security Type drop-down list, select WPA-Enterprise or WPA2-Enterprise. In Choose A Network Authentication Method, select Protected EAP (PEAP), and then click Settings.
3. In the Protected EAP Properties dialog box, if you want to validate the computer certificate of the NPS server for the PEAP authentication, select Validate Server Certificate (recommended and enabled by default). If you want to specify the names of the NPS servers that must perform the PEAP authentication, select Connect To These Servers, and then type the names.
4. In the Select Authentication Method drop-down list, click Smart Card Or Other Certificate. Click Configure. To use a registry-based user certificate, in the Smart Card Or Other Certificate Properties dialog box, select Use A Certificate On This Computer. For a smart card–based user certificate, select Use My Smart Card.

   If you want to validate the computer certificate of the NPS server for the user-level authentication, select the Validate Server Certificate check box (recommended and enabled by default). If you want to specify the names of the NPS servers that must perform the TLS authentication, select Connect To These Servers, and then type the names.

5. Click OK to save changes to the Smart Card or other Certificate PEAP type. Click OK to save the changes to the Protected EAP type. Click OK to save the changes to the wireless network configuration.

To manually configure PEAP-TLS authentication on a wireless client running Windows XP with SP2, Windows XP with SP1, or Windows Server 2003, do the following:

1. Obtain properties of the wireless connection in the Network Connections folder. On the Wireless Networks tab, in the list of preferred networks, click the name of the wireless network, and then click Properties. The Wireless Network's properties dialog box appears.

2. On the Authentication tab, select Enable Network Access Control Using IEEE 802.1X and the Protected EAP (PEAP) type.
3. Click Properties. In the Protected EAP Properties dialog box, select the Validate Server Certificate check box to validate the computer certificate of the NPS server for the PEAP authentication (recommended and enabled by default). If you want to specify the names of the authentication servers that must perform PEAP authentication, select Connect To These Servers, and then type the names. In the Select Authentication Method drop-down list, click Smart Card Or Other Certificate.
4. Click Configure. In the Smart Card Or Other Certificate Properties dialog box, to use a registry-based user certificate, select Use A Certificate On This Computer. For a smart card–based user certificate, select Use My Smart Card.

   If you want to validate the computer certificate of the NPS server for the user-level authentication, select Validate Server Certificate (recommended and enabled by default). If you want to specify the names of the NPS servers that must perform the TLS authentication, select Connect To These Servers, and then type the names.
5. Click OK to save changes to the Smart Card or other Certificate PEAP type. Click OK to save the changes to the Protected EAP type. Click OK to save the changes to the wireless network configuration.

PEAP-MS-CHAP v2 To manually configure PEAP-MS-CHAP v2 authentication on a wireless client running Windows Server 2008 or Windows Vista, do the following:

1. In the Network and Sharing Center, click the Manage Wireless Networks task. In the Manage Wireless Networks window, double-click your wireless network name.
2. On the Security tab, in the Security Type drop-down list, select WPA-Enterprise or WPA2-Enterprise. In the Choose a network authentication method drop-down list, select Protected EAP (PEAP), and then click Settings.
3. In the Protected EAP Properties dialog box, if you want to validate the computer certificate of the NPS server for the PEAP authentication, select the Validate Server Certificate check box (recommended and enabled by default). If you want to specify the names of the NPS servers that must perform the PEAP authentication, select Connect To These Servers, and then type the names.
4. In Select Authentication Method, select Secured Password (EAP-MS-CHAP v2), and then click OK twice.

To manually configure PEAP-MS-CHAP v2 authentication on a wireless client running Windows XP with SP2, Windows XP with SP1, or Windows Server 2003, do the following:

1. Obtain properties of the wireless connection in the Network Connections folder. Click the Wireless Networks tab, click the name of the wireless network in the list of preferred networks, and then click Properties. The wireless network's properties dialog box appears.

2. On the Authentication tab, select Enable Network Access Control Using IEEE 802.1X and the Protected EAP (PEAP) EAP type.
3. Click Properties. In the Protected EAP Properties dialog box, select Validate Server Certificate to validate the computer certificate of the NPS server (enabled by default). If you want to specify the names of the authentication servers that must perform validation, select Connect To These Servers, and then type the names. In Select Authentication Method, click Secured Password (EAP-MSCHAP v2), and then click OK twice.

Ongoing Maintenance

The areas of maintenance for a protected wireless solution are as follows:

- Management of user and computer accounts
- Management of wireless APs
- Updating of wireless profiles

Managing User and Computer Accounts

When a new user or computer account is created in Active Directory, and that user or computer is allowed wireless access, add the new account to the appropriate group for wireless connections. For example, add the new account to the WirelessAccounts security group, which is specified in the network policy for wireless connections.

When user or computer accounts are deleted in Active Directory, no additional action is necessary to prevent wireless connections.

As needed, you can create additional universal groups and network policies to set wireless network access for different sets of users. For example, you can create a global WirelessAccessContractors group and a network policy that allows wireless connections to members of the WirelessAccessContractors group only during normal business hours or for access to specific intranet resources.

Managing Wireless APs

Once deployed, wireless APs do not need a lot of ongoing maintenance. Most of the ongoing changes to wireless AP configuration are due to managing wireless network capacity and changes in network infrastructure.

Adding a Wireless AP

To add a wireless AP, do the following:

1. Follow the design points and deployment steps in “Deploying Wireless APs” earlier in this chapter to add a new wireless AP to your wireless network.
2. Add the wireless AP as a RADIUS client to your NPS servers.

Removing a Wireless AP

When removing a wireless AP, update the configuration of your NPS servers to remove the wireless AP as a RADIUS client.

Configuration for Changes in NPS Servers

If the NPS servers change (for example, because of additions or removals of NPS servers on the intranet), you will need to do the following:

1. Ensure that new NPS servers are configured with RADIUS clients corresponding to the wireless APs and with the appropriate network policies for wireless access.
2. Update the configuration of the wireless APs for the new NPS server configuration as needed.

Updating Wireless XML Profiles

To update a wireless XML profile and apply it to your Windows Vista or Windows Server 2008 wireless clients, do the following:

1. If you are using a Windows Vista or Windows Server 2008 wireless client or if you have a Windows Vista wireless policy, create an updated XML profile with the Group Policy Editor snap-in or by running the **netsh wlan export profile** command.
2. Execute the **netsh wlan add profile** command to import the XML profile on your wireless clients through a script or other method.

Troubleshooting

Because of the different components and processes involved, troubleshooting wireless connections can be a difficult task. This section describes the following:

- The tools that are provided with Windows Server 2008 and Windows Vista to troubleshoot wireless connections
- How to troubleshoot wireless connection problems from the wireless client
- How to troubleshoot wireless connection problems from the wireless AP
- How to troubleshoot wireless connection problems from the NPS server

Direct from the Source: Wireless Troubleshooting Tips

One of the most difficult aspects of troubleshooting wireless connectivity is knowing where to start. Generally, the client is the device that shows the symptom, but it is only one piece in a chain of devices and technologies that could fail.

As a general rule to follow, if the wireless client fails to see the wireless network or establish an association, the issue lies between the wireless client and the wireless AP. Most of these issues are resolved by driver or firmware updates for the wireless network adapter and the wireless AP. Having the latest drivers and firmware installed is a required first step in the troubleshooting process.

If authentication is failing, you most likely can rule out hardware as an issue. First review your client-side System event logs. Windows XP and Windows Server 2003 do not have any diagnostic logs, but Windows Server 2008 and Windows Vista log quite a bit of useful information that might point you to a configuration issue such as a missing certificate.

After reviewing these logs, review the Windows Logs\Security event log on the NPS server. If you have a failed authentication, there will be an NPS event with the keyword Audit Failure. If, however, you do not see any log entries related to the wireless authentication attempt, this is a strong indicator that NPS did not receive the authentication attempt or the process timed out. Take a look at the wireless AP to confirm that its RADIUS settings are appropriate for the NPS server.

Clay Seymour, Support Escalation Engineer

Enterprise Platform Support

Wireless Troubleshooting Tools in Windows

Microsoft provides the following tools to troubleshoot wireless connections:

- TCP/IP troubleshooting tools
- The Network Connections folder
- **Netsh wlan** commands
- Network Diagnostics Framework support for wireless connections
- Wireless diagnostics tracing
- NPS authentication and accounting logging
- NPS event logging
- SChannel logging
- SNMP agent
- Reliability and Performance snap-in
- Network Monitor 3.1

TCP/IP Troubleshooting Tools

The Ping, Tracert, and Pathping tools use Internet Control Message Protocol (ICMP) Echo and Echo Reply and ICMPv6 Echo Request and Echo Reply messages to verify connectivity, display the path to a destination, and test path integrity. The Route tool can be used to display the IPv4 and IPv6 routing tables. The Nslookup tool can be used to troubleshoot domain name system (DNS) name resolution issues.

The Network Connections Folder

When you obtain status on the wireless connection in the Network Connection folder, you can view information such as the signal speed, which is shown on the General tab. Click Details to view the TCP/IP configuration.

If the wireless adapter is assigned an Automatic Private IP Addressing (APIPA) address in the range 169.254.0.0/16 or the configured alternate IP address, the wireless client is still associated with the wireless AP, but either authentication has failed or the DHCP server is not available. If the authentication fails and the association is still in place, the wireless adapter is enabled and TCP/IP performs its normal configuration process. If a DHCP server is not found (either authenticated or not), Windows Vista automatically configures an APIPA address unless there is an alternate address configured.

Direct from the Source: APIPA in Windows Vista

You might notice that a Windows Vista wireless client will automatically configure an APIPA address sooner or more frequently than in previous versions of Windows. A computer running Windows Vista will wait only six seconds to contact a DHCP server before using an APIPA address and will then continue to attempt to contact a DHCP server. By contrast, a computer running Windows XP will wait a full minute before using an APIPA address. This change in behavior is by design and is meant to facilitate ad-hoc connectivity when there are no DHCP servers available.

Tim Quinn, Support Escalation Engineer

Enterprise Platform Support

Netsh Wlan Commands

You can run the **netsh wlan** command with the following parameters to gather information for troubleshooting wireless issues:

- **netsh wlan show autoconfig** Displays whether the WLAN Autoconfig service is enabled
- **netsh wlan show blockednetworks** Displays whether blocked networks are visible in the list of available networks

- **netsh wlan show createalluserprofile** Displays whether everyone is allowed to create all-user profiles
- **netsh wlan show drivers** Displays the properties of the drivers for the installed wireless network adapters
- **netsh wlan show filters** Displays the allowed and blocked wireless networks lists
- **netsh wlan show interfaces** Displays properties for the installed wireless network adapters
- **netsh wlan show networks** Displays the list and properties of the available wireless networks
- **netsh wlan show profiles** Displays the list of Group Policy and local wireless profiles
- **netsh wlan show settings** Displays the global wireless settings, which includes the state of Wireless Auto Configuration and whether everyone is allowed to create all-user profiles.
- **netsh wlan show tracing** Displays the state of tracing and the location of the wireless tracing logs (by default in *%SystemRoot%*\Tracing\Wireless)
- **netsh wlan show all** Displays complete wireless network adapter information and information on available wireless networks

Network Diagnostics Framework Support for Wireless Connections

To provide a better user experience when encountering network connectivity issues, Windows Vista includes the Network Diagnostics Framework (NDF), a set of technologies and guidelines that allows a set of troubleshooters (also known as *helper classes*) to assist in the diagnosis and possible automatic correction of networking problems. When a user experiences a networking problem in Windows Vista, NDF will provide the user the ability to diagnose and repair the problem within the context of that problem. This means that the diagnostics assessment and resolution steps are presented to the user within the application or dialog box that they were using when the problem occurred or based on the failed network operation.

Windows Vista includes a troubleshooter to diagnose failed wireless connections. If a wireless connection fails, Windows displays a dialog box with information about the error. The dialog box includes a Diagnose button that launches the wireless NDF troubleshooter. In the diagnosis session, users can repair their wireless connection problem without needing to involve IT support staff. The wireless NDF troubleshooter will help users resolve many common issues that arise with wireless network connectivity, such as:

- The network adapter radio being turned off
- The wireless AP not being powered

- A missing or mismatched configuration of security options, encryption types, or network keys between the wireless AP and wireless client
- Disconnected media
- Missing certificates

Windows logs all wireless connection attempts in the System event log. When Windows Network Diagnostics runs, it creates additional events in the System event log that contain the following information:

- The name of the wireless network adapter and whether its driver is designed for Windows Vista.
- A list of visible wireless networks with the signal strength, channel, protocol (such as 802.11b or 802.11g), and operating mode (infrastructure or ad hoc) for each.
- The list of preferred wireless networks and each network's configuration settings.
- The diagnostic conclusions, such as, "The wireless connection on this computer appears to be working correctly," "The Internet connection on the wireless router or access point might not be working correctly," and "The computer has a low signal strength from ContosoWLAN."
- The repair options offered to the user, such as, "Try moving the computer to a different location, eliminating any sources of possible interference, and then try connecting to ContosoWLAN again."
- The repair options chosen by the user and whether the repair solved the problem.

You can view these events in the Event Viewer snap-in to understand the network environment at the time the problem occurred without needing to re-create the scenario, and you need no longer rely on users to explain the symptoms of the problem.

To obtain additional information about the diagnostics process, Windows creates a detailed diagnostic log that is separate from the System event log.

To Access The Diagnostics Log

1. In the Event Viewer snap-in, in the tree view, expand Applications and Services Logs\Microsoft\Windows\Diagnostics-Networking.
2. Click Operational.
3. In the contents pane, view the events for the wireless diagnostics session.

Wireless Diagnostics Tracing

Occasionally, you might need to escalate a wireless networking problem to Microsoft or another support specialist in your organization. To perform a detailed analysis, Microsoft or your support specialists need in-depth information about the computer's state and wireless components in Windows and their interaction when the problem occurred. You can obtain this information from wireless diagnostics tracing in Windows Vista. To use wireless diagnostics tracing, you must start tracing, reproduce the problem, stop tracing, and then collect the tracing report.

To start wireless diagnostics tracing, do one of the following:

- Type the **netsh wlan set tracing mode=yes** command at a command prompt.
- In the console tree of the Reliability and Performance Monitor snap-in, expand Data Collector Sets\System. Right-click Wireless Diagnostics, and then click Start.

After you have reproduced the problem and want to stop wireless diagnostics tracing, do one of the following:

- Type the **netsh wlan set tracing mode=no** command.
- In the console tree of the Reliability and Performance Monitor snap-in, expand Data Collector Sets\System. Right-click Wireless Diagnostics, and then click Stop.

> **Note** It is important to stop the wireless diagnostics tracing prior to viewing or gathering the trace logs to initiate a process that converts the trace files into a readable format.

To view the report generated by wireless diagnostics tracing, in the console tree of the Reliability and Performance Monitor snap-in, expand Reports\System\Wireless Diagnostics.

The report includes the following information:

- Wireless configuration, including allowed and blocked wireless networks
- Current TCP/IP configuration (including data provided by the **ipconfig /all** command)
- A list of all connection attempts and detailed information about each step of the connection process
- A detailed list of all Windows Network Diagnostics events
- Wireless client certificate configuration
- Wireless profiles and their locations
- Wireless network adapter driver information
- Wireless networking system files and versions
- Raw network tracing information
- Computer make and model
- Operating system version
- A list of all services, their current states, and their process identifiers

This report and its associated files are stored by default in the *%SystemRoot%*\Tracing\Wireless folder.

In addition to wireless diagnostic tracing, Windows Server 2008 and Windows Vista support tracing for components of the Remote Access Connection Manager and Routing and Remote Access services, which are also used for wireless connections. Like the wireless diagnostic tracing, tracing for these components creates information that you can use to troubleshoot complex problems for specific components. The information in these additional tracing files is typically useful only to Microsoft support engineers, who might request that you create trace files for a connection attempt during their investigation of a support issue. You can enable this additional tracing by using the Netsh tool.

To enable and disable tracing for a specific component of the Remote Access Connection Manager and Routing and Remote Access services, the command is:

netsh ras diagnostics set rastracing *component* enabled|disabled

in which ***component*** is a component in the list of components found in the registry under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Tracing.

To enable tracing for all components, the command is:

netsh ras diagnostics set rastracing * enabled

To disable tracing for all components, the command is:

netsh ras diagnostics set rastracing * disabled

The tracing log files are stored in the *%SystemRoot%*\Tracing folder. The most interesting log files for wireless authentication are the following:

- **Svchost_rastls.log** TLS authentication activity
- **Svchost_raschap.log** MS-CHAP v2 authentication activity

NPS Authentication and Accounting Logging

By default, NPS supports the logging of authentication and accounting information for wireless connections in local log files. This logging is separate from the events recorded in the Windows Logs\Security. You can use the information in the logs to track wireless usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting network policy issues. For each authentication attempt, the name of the network policy that either accepted or rejected the connection attempt is recorded. You can configure authentication and accounting logging options in the Accounting node in the Network Policy Server snap-in.

The authentication and accounting information is stored in a configurable log file or files stored in the *%SystemRoot%*\System32\LogFiles folder. The log files are saved in Internet

Authentication Service (IAS) or database-compatible format, meaning that any database program can read the log file directly for analysis. NPS can also send authentication and accounting information to a SQL Server database.

NPS Event Logging

Check the Windows Logs\Security event log on the NPS server for NPS events corresponding to rejected (event ID 6273) or accepted (event ID 6272) connection attempts. NPS event log entries contain a lot of information on the connection attempt, including the name of the connection request policy that matched the connection attempt (the Proxy Policy Name in the description of the event) and the network policy that accepted or rejected the connection attempt (the Network Policy Name field in the description of the event). NPS event logging for rejected or accepted connection attempts is enabled by default. You can configure it in the Network Policy Server snap-in, in the properties dialog box of an NPS server, on the General tab.

NPS events can be viewed from the Event Viewer snap-in. Viewing the NPS events in the Windows Logs\Security event log is one of the most useful troubleshooting methods to obtain information about failed authentications.

SChannel Logging

Secure channel (SChannel) logging is the logging of detailed information for SChannel events in the System event log. By default, only SChannel error messages are recorded. To log errors, warnings, and informational and successful events, set the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurityProviders\SCHANNEL\EventLogging registry value to **4** (DWORD value type). With SChannel logging recording all events, it is possible to obtain more information about the certificate exchange and validation process on the NPS server.

SNMP Agent

You can use the Simple Network Management Protocol (SNMP) agent software included with Windows Server 2008 to monitor status information for your NPS server from an SNMP console. NPS supports the RADIUS Authentication Server MIB (RFC 2619) and the RADIUS Accounting Server MIB (RFC 2621). Use Features in the Server Manager console to install the optional SNMP service.

The SNMP service can be used in conjunction with your existing SNMP-based network management infrastructure to monitor your NPS RADIUS servers or proxies.

Reliability and Performance Snap-In

You can use the Reliability and Performance snap-in to monitor counters, create logs, and set alerts for specific NPS components and program processes. You can also use charts and reports to determine how efficiently your server uses NPS and to both identify and troubleshoot potential problems.

You can use the Reliability and Performance snap-in to monitor counters within the following NPS-related performance objects:

- NPS Accounting Clients
- NPS Accounting Proxy
- NPS Accounting Server
- NPS Authentication Clients
- NPS Authentication Proxy
- NPS Authentication Server
- NPS Policy Engine
- NPS Remote Accounting Servers
- NPS Remote Authentication Servers

More Info For more information about how to use the Reliability and Performance snap-in, see the Help And Support Center in Windows Server 2008.

Network Monitor 3.1

You can use Microsoft Network Monitor 3.1 (or later) or a commercial packet analyzer (also known as a network sniffer), to capture and view the authentication and data traffic sent on a network. Network Monitor 3.1 includes RADIUS, 802.1X, EAPOL, and EAP parsers. A *parser* is a component included with Network Monitor that can separate the fields of a protocol header and display their structure and values. Without a parser, Network Monitor 3.1 displays the hexadecimal bytes of a header, which you must parse manually.

On the Disc You can link to the download site for Network Monitor from the companion CD-ROM.

For Windows wireless client authentications, you can use Network Monitor 3.1 to capture the set of frames exchanged between the wireless client computer and the wireless AP during the wireless authentication process. You can then use Network Monitor 3.1 to view the individual frames and determine why the authentication failed. Network Monitor 3.1 is also useful for capturing the RADIUS messages that are exchanged between a wireless AP and its RADIUS server and for determining the RADIUS attributes of each message.

The proper interpretation of wireless traffic with Network Monitor 3.1 requires an in-depth understanding of EAPOL, RADIUS, and other protocols. Network Monitor 3.1 captures can be saved as files and sent to Microsoft support for analysis.

Troubleshooting the Windows Wireless Client

When troubleshooting wireless connectivity, it is important to first determine whether some or all of your wireless clients are experiencing problems. If all your wireless clients are experiencing problems, issues might exist in your authentication infrastructure. If some of your wireless clients are experiencing problems, issues might exist for your wireless APs or individual wireless clients.

The following are some common problems with wireless connectivity and authentication that are encountered by a Windows wireless client:

- **Wireless network is not found.** Verify that you are within range of the wireless AP for the wireless network by using tools provided by the wireless adapter vendor. You can move the wireless AP or the wireless client, adjust the transmission power level on the wireless AP, or reposition or remove sources of radio frequency attenuation or interference.
- **Unable to authenticate.** Some wireless network adapters have a link light that indicates sent or received data frames. However, because IEEE 802.1X authentication occurs before the wireless network adapter begins sending or receiving data frames, the link light does not reflect 802.1X authentication activity. If the link light does not indicate any wireless traffic, the cause could be a failed 802.1X authentication.

  Verify that the user or computer account for the wireless client exists, is enabled, and is not locked out (via account properties or remote access account lockout), and that the connection is being attempted during allowed logon times.

  Verify that the connection attempt for the user or computer account matches a network policy. For example, if you are using a group-based network policy, verify that the user or computer account is a member of the group specified in the Windows Groups condition of the appropriate network policy.

  Verify that the root CA certificates for the issuing CAs of the NPS server certificates are installed in the Trusted Root Certification Authorities Local Computer store on the wireless client computer.

  For an EAP-TLS–based or PEAP-TLS–based wireless client, verify that the computer or user certificate meets the conditions described in the section titled “Validating the Wireless Client’s Certificate” later in this chapter.

  For a PEAP-MS-CHAP v2–based wireless client, investigate whether the wireless client’s account password has expired, and verify that the Allow Client to Change Password After It Has Expired check box in the EAP MS-CHAP v2 Properties dialog box is enabled on the NPS servers.
- **Unable to authenticate with a certificate.** The most typical cause for this message is that you do not have either a user or computer certificate installed. Depending on the

configured authentication mode, you might need to have both installed. Verify that you have a computer certificate, a user certificate, or both installed by using the Certificates snap-in.

Another possible cause for this message is that you have certificates installed, but they either cannot be used for wireless authentication, or they cannot be validated by all of your NPS servers. For more information, see "Troubleshooting Certificate-Based Validation" later in this chapter.

Troubleshooting the Wireless AP

If you have multiple wireless APs and are unable to connect or authenticate with one of them, you might have a problem with that specific wireless AP. This section describes the common troubleshooting tools of wireless APs and the common problems of connecting and authenticating with a wireless AP.

Wireless AP Troubleshooting Tools

Although the set of troubleshooting tools for wireless APs varies with each manufacturer and with each model, some of the more common troubleshooting tools are the following:

- Panel indicators
- Site survey software
- SNMP support
- Diagnostics

These tools are described in the following sections. Consult your wireless AP documentation for information about the set of troubleshooting tools provided with your wireless AP.

Panel Indicators Most wireless APs have one or more indicators, which are status lights that are visible on the housing of the wireless AP, from which you can obtain a quick assessment of the wireless AP's hardware status. For example, you might see the following:

- An indicator to show that the wireless AP has electrical power.
- An indicator to show general operation status. For example, the indicator might show whether the wireless AP is associated with any wireless clients.
- An indicator to show wireless network traffic. This indicator might blink for each frame received on the wireless network.
- An indicator to show data collisions. If the blinking of this indicator seems excessive, evaluate the performance of the link by using the methods suggested by the wireless AP vendor.
- An indicator to show wired network traffic. This indicator might blink for each frame received on the wired network.

Alternatively, the wireless AP might have a liquid crystal display (LCD) panel that shows icons that indicate its current status. Consult your wireless AP documentation for information about panel indicators and their interpretation.

Site Survey Software Site survey software, which you use during the deployment of wireless APs to determine their optimal placement, is typically installed on a wireless-capable laptop computer from a CD-ROM provided by the wireless AP or wireless network adapter vendor.

As described in "Deploying Wireless APs," earlier in this chapter, the site survey software is used to determine the coverage volume and where the data rate changes for each wireless AP. If wireless clients cannot connect to a specific wireless AP, use the site survey software to perform a site survey for that wireless AP. There might have been a change in the devices that create interference and objects that interfere with signal propagation since the original site survey and AP placement were done.

SNMP Support Many wireless APs include a Simple Network Management Protocol (SNMP) agent with support for the following SNMP Management Information Bases (MIBs):

- IEEE 802.11 MIB
- IEEE 802.1 PAE (Port Access Entity) MIB
- SNMP Management MIB (described in RFC 1157)
- SNMP MIB II (described in RFC 1213)
- Bridge MIB (described in RFC 1286)
- Ethernet Interface MIB (described in RFC 1398)
- IETF Bridge MIB (described in RFC 1493)
- Remote Monitoring (RMON) MIB (described in RFC 1757)
- RADIUS Client Authentication MIB (described in RFC 2618)

The SNMP agent on the wireless AP can be used in conjunction with your existing SNMP-based network management infrastructure to configure your wireless APs, set trap conditions, and monitor loads on your wireless APs.

Diagnostics Diagnostics for wireless APs can be in the following forms:

- Diagnostic facilities that are available through the main wireless AP configuration program, such as a Windows program provided on the wireless AP vendor product CD-ROM or a series of Web pages.
- Diagnostic facilities that are available through a command-line tool or facility, such as terminal access to the wireless AP.

The exact diagnostic facilities of a wireless AP vary from one wireless AP to another; however, the purpose of the diagnostics is to ensure that the wireless AP is operating properly (from a hardware standpoint) and to validate its current configuration.

Common Wireless AP Problems

The following are common problems with wireless APs:

- Inability to see the wireless AP
- Inability to authenticate with the wireless AP
- Inability to communicate beyond the wireless AP

These common problems are discussed in detail in the following sections.

Inability to See the Wireless AP If wireless clients are unable to see the wireless AP in a scan of wireless networks, one or more of the following might be happening:

- **The wireless AP is not beaconing.** All wireless APs should be sending periodic beacon messages that contain the SSID—unless the wireless AP has been configured to suppress the SSID in the beacon message—and the wireless AP's capabilities (such as supported bit rates and security options). To verify that the wireless AP is beaconing, you can use the site survey software or a packet sniffer that can capture wireless beacon frames. A simple packet sniffer that can capture beacon frames and other types of wireless management frames might be included on the CD-ROM provided by your wireless AP vendor.
- **The wireless AP is not configured for the correct channel.** If the wireless AP is using the same channel as an adjacent wireless AP, signal interference might be impairing the wireless clients' ability to connect. Change the wireless AP channel if needed.
- **The wireless AP is not advertising the correct set of capabilities.** Confirm that the wireless AP is configured to operate for the correct technology (such as 802.11b, 802.11a, or 802.11g) and with the correct bit rates and security options (WPA or WPA2). By capturing the beacon frame with a network sniffer, you can compare the configured wireless options to those being advertised in the beacon frame.
- **The wireless AP has inadequate signal strength in the anticipated coverage volume.** Use your site survey software to confirm that the coverage volume of the wireless AP is as described in your plans after initially deploying the wireless APs. If there are new sources of signal attenuation, reflection, or interference, make the appropriate changes to the locations of either interfering equipment or the wireless AP.

Inability to Authenticate with the Wireless AP If you have multiple wireless APs, and your wireless clients cannot authenticate with any of them, you might have a problem with your authentication infrastructure. See "Troubleshooting the Authentication Infrastructure" later in this chapter for instructions on how to troubleshoot this situation. If you have multiple wireless APs, and the wireless clients cannot authenticate with an individual wireless AP, you need to troubleshoot the authentication-related configuration of the wireless AP. The three areas of authentication configuration you need to investigate are as follows:

- 802.1X configuration

- RADIUS configuration
- WPA configuration

802.1X Configuration Ensure that the wireless AP has 802.1X authentication enabled. Some wireless APs might refer to 802.1X authentication as EAP authentication.

RADIUS Configuration The RADIUS configuration consists of the following elements:

- **Wireless AP RADIUS configuration** Ensure that the wireless AP has been properly configured for RADIUS. The wireless AP should contain the following configuration information:
  - The IPv4 or IPv6 address of a primary RADIUS server (one of your NPS servers)
  - The destination User Datagram Protocol (UDP) ports for RADIUS traffic sent to the primary RADIUS server (UDP port 1812 for RADIUS authentication traffic and UDP port 1813 for RADIUS accounting traffic)
  - The RADIUS shared secret for the primary RADIUS server
  - The IPv4 or IPv6 address of a secondary RADIUS server (another of your NPS servers)
  - The destination UDP ports for RADIUS traffic sent to the secondary RADIUS server
  - The RADIUS shared secret for the secondary RADIUS server
- **NPS server reachability** Ensure that the primary and secondary NPS servers are reachable from the wireless AP by doing the following:
  - If the wireless AP has a ping facility—the capability to send an Internet Control Message Protocol (ICMP) Echo message to an arbitrary unicast IPv4 destination—try pinging the IPv4 address of the primary and secondary NPS servers.
  - If the wireless AP does not have a ping facility, try pinging the IPv4 address of the primary and secondary NPS servers from a network node that is attached to the same subnet as the wireless AP.

  If the ping from the network node succeeds and the ping from the wireless AP does not, examine the IPv4 configuration of the wireless AP to ensure that it has been configured with the correct IPv4 address, subnet mask, and default gateway for the attached wired subnet. If neither ping works, troubleshoot the lack of IPv4 connectivity between the attached subnet and the RADIUS servers.

> **Note** The ping test is not necessarily a definitive test of IPv4 reachability. There might be routers in the path between the wireless AP and the RADIUS server that are filtering ICMP traffic, or the NPS server might be configured with packet filters to discard ICMP traffic.

To ensure that RADIUS traffic is reaching the primary and secondary NPS servers, use a network sniffer such as Network Monitor 3.1 on the NPS servers to capture the RADIUS traffic sent from and to the wireless AP during an authentication attempt.

- **NPS server configuration** If RADIUS traffic is reaching the primary and secondary NPS servers, verify that the primary and secondary NPS servers are configured with a RADIUS client that corresponds to the wireless AP, including the following:
    - The IPv4 address of the wireless AP's wired interface
    - The destination UDP ports for RADIUS traffic sent by the wireless AP (UDP port 1812 for RADIUS authentication traffic and UDP port 1813 for RADIUS accounting traffic)
    - The RADIUS shared secret configured at the wireless AP

  Check the Windows Logs\Security event log for authentication failure events corresponding to connection attempts to the wireless AP. To view the failed authentication events, use the Event Viewer to view the events in the Security event log with the event ID of 6273.
- **IPsec for RADIUS traffic** If you are using IPsec to encrypt the RADIUS traffic sent between the wireless AP and the NPS server, check the IPsec settings on both the wireless AP and NPS server to ensure that they can successfully negotiate security associations and authenticate each other.

> **Note** For more information about how to configure IPsec policies in Windows Server 2008 to provide protection for RADIUS traffic, see Chapter 4, "Windows Firewall with Advanced Security." For more information about how to configure IPsec settings for a wireless AP, see your wireless AP's product documentation.

WPA or WPA2 Configuration If your wireless AP is WPA-capable or WPA2-capable and you want to use WPA or WPA2 for wireless security, ensure that WPA or WPA2 is enabled.

Inability to Communicate Beyond the Wireless AP The wireless AP is a transparent bridge and Layer 2 switching device, forwarding packets between the wired network to which it is attached and the connected wireless clients. If wireless clients can connect and authenticate but cannot reach locations beyond the wireless AP, one or more of the following might be happening.

- **The wireless AP is not forwarding frames as a bridge.** All transparent bridges support the spanning tree protocol, which is used to prevent loops in a bridged section of the network. The spanning tree protocol uses a series of multicast messages to communicate bridge configuration information and automatically configure bridge interfaces to forward frames or block forwarding to prevent loops. While the spanning tree algorithm is determining forwarding and blocking interfaces, the bridge is not forwarding frames. Check the wireless AP's forwarding status and bridge configuration.

- **The wireless AP is not configured with the correct VLAN IDs.** Many wireless APs support VLANs, which are switch ports grouped so that they appear on the same link or subnet. Each group is assigned a separate VLAN ID. Verify that the VLAN IDs for your wireless client and your wired interfaces are correctly configured. For example, you might use one VLAN ID for authenticated wireless clients (that connects them to the organization intranet) and a separate VLAN ID for guest wireless clients (that connects them to an alternate subnet or the Internet).

Troubleshooting the Authentication Infrastructure

If you have multiple wireless APs and are unable to authenticate with any of them, you might have a problem with your authentication infrastructure, which consists of your NPS servers, PKI, and Active Directory accounts. In this section we examine common issues with NPS authentication and authorization, and validation of certificate-based and password-based authentications.

Troubleshooting NPS Authentication and Authorization

To troubleshoot the most common issues with NPS authentication and authorization, verify the following:

- **That the wireless AP can reach the NPS servers** To test this, try to ping the IP address of the wireless AP's interface on the wired network from each of the NPS servers. Additionally, ensure that IPsec policies, IP packet filters, and other mechanisms that restrict network traffic are not preventing the exchange of RADIUS messages between the wireless AP and its configured NPS servers. RADIUS traffic to the NPS servers uses a source IPv4 or IPv6 address of the wireless AP, a destination IPv4 or IPv6 address of the NPS server, and a UDP destination port of 1812 for authentication messages and UDP destination port 1813 for accounting messages. RADIUS traffic from the NPS servers uses a source IPv4 or IPv6 address of the NPS server, a destination IPv4 or IPv6 address of the wireless AP, a UDP source port of 1812 for authentication messages, and UDP source port 1813 for accounting messages. These examples assume that you are using the RADIUS UDP ports defined in RFC 2865 and 2866 for RADIUS authentication and accounting traffic.
- **That each NPS server/wireless AP pair is configured with a common RADIUS shared secret** Each NPS server/wireless AP pair is not necessarily required to use a unique RADIUS shared secret, but it must use the same value for the RADIUS shared secret for the members of the pair. For example, when you copy the NPS configuration from one NPS server to another, verify all of the shared secret pairs between the NPS servers and the wireless APs.
- **That the NPS servers can reach a global catalog server and an Active Directory domain controller** The NPS server uses a global catalog server to resolve the user principal name (UPN) of the computer or user certificate or the MS-CHAP v2 account

name to the distinguished name of the corresponding account in Active Directory. The NPS server uses an Active Directory domain controller to validate the credentials of the computer and user account and obtain account properties to evaluate authorization.

- **That the computer accounts of the NPS servers are members of the RAS and IAS Servers security group for the appropriate domains** Adding the NPS server computer accounts to the RAS and IAS Servers security group for the appropriate domains is normally done during the initial configuration of the NPS server. To add the NPS server computer account to the appropriate domains, you can run the **netsh nps add registeredserver** command.
- **That there are no configured restrictions blocking access** Ensure that the user or computer account is not locked out, expired, or disabled or that the time the connection is being made corresponds to the permitted logon hours.
- **That the user account has not been locked out by remote access account lockout** Remote access account lockout is an authentication counting and lockout mechanism designed to prevent an online dictionary attack against a user's password. If remote access account lockout is enabled, you can reset account lockout for the account by deleting the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RemoteAccess\Parameters\AccountLockout\DomainName:*AccountName* registry value on the NPS server.
- **That the connection is authorized** For authorization, the parameters of the connection attempt must:
  - ❑ Match all the conditions of at least one network policy. If there is no matching policy, all wireless authentication requests are rejected.
  - ❑ Be granted network access permission through the user account (set to Allow Access), or if the user account has the Control Access Through NPS Network Policy option selected, the access permission of the first matching network policy must be set to Grant Access.
  - ❑ Match all the settings of the profile. Verify that the authentication settings of the profile have EAP-TLS or PEAP-MS-CHAP v2 enabled and properly configured.
  - ❑ Match all the settings of the dial-in properties of the user or computer account.

  To obtain the name of the network policy that rejected the connection attempt, ensure that NPS event logging is enabled for rejected authentication attempts, and use the Event Viewer to view the events in the Windows Logs\Security event log that have the event ID of 6273. In the text of the event for the connection attempt, look for the network policy name in the Network Policy Name field.
- **That you have not changed the mode of your domain from mixed mode to native mode** If you have just changed your Active Directory domain from mixed mode to native mode, NPS servers can no longer authenticate valid connection requests. You must restart every domain controller in the domain for the change to replicate.

Troubleshooting Certificate-Based Validation

Troubleshooting certificate validation for EAP-TLS or PEAP-TLS authentication consists of verifying the wireless client's computer and user certificates and the computer certificates of the NPS servers.

Validating the Wireless Client's Certificate For an NPS server to validate the certificate of a wireless client, the following must be true for each certificate in the certificate chain sent by the wireless client:

- **The current date is within the validity dates of the certificate.** When certificates are issued, they are issued with a valid date range, before which they cannot be used and after which they are considered expired.
- **The certificate has not been revoked.** Issued certificates can be revoked at any time. Each issuing CA maintains a list of certificates that should no longer be considered valid by publishing an up-to-date certificate revocation list (CRL). The server will first attempt to validate the certificate using the Online Certificate Status Protocol (OSCP). If the OSCP validation is successful, the validation verification is satisfied; otherwise, it will then attempt to perform a CRL validation of the user or computer certificate. By default, the NPS server checks all the certificates in the wireless client's certificate chain (the series of certificates from the wireless client certificate to the root CA) for revocation. If any of the certificates in the chain have been revoked, certificate validation fails. This behavior can be modified by changing registry settings as described later in this chapter.

  To view the CRL distribution points for a certificate in the Certificates snap-in, in the contents pane, double-click the certificate, click the Details tab, and then click the CRL Distribution Points field. To perform a revocation check, the NPS server must be able to reach the CRL distribution points.

  The certificate revocation check works only as well as the CRL publishing and distribution system. If the CRL is not updated often, a certificate that has been revoked can still be used and considered valid because the published CRL that the NPS server is checking is out of date. Verify that the CRLs available to the NPS servers have not expired. If the CRLs available to the NPS servers have expired, EAP-TLS and PEAP-TLS authentication fails.
- **The certificate has a valid digital signature.** CAs digitally sign certificates they issue. The NPS server verifies the digital signature of each certificate in the chain (with the exception of the root CA certificate) by obtaining the public key from the certificate's issuing CA and mathematically validating the digital signature.

  The wireless client certificate must also have the Client Authentication certificate purpose (also known as Enhanced Key Usage, or EKU) and must contain either a UPN of a valid user account or a Fully Qualified Domain Name (FQDN) of a valid computer account in the Subject Alternative Name field of the certificate.

  To view the EKU for a certificate in the Certificates snap-in, double-click the certificate in the contents pane, and then on the Details tab, click the Enhanced Key Usage field.

To view the Subject Alternative Name field for a certificate in the Certificates snap-in, in the contents pane, double-click the certificate, click the Details tab, and then click the Subject Alternative Name field.

- **The NPS server must have the appropriate certificate installed correctly.** To trust the certificate chain offered by the wireless client, the NPS server must have the root CA certificate of the issuing CA of the wireless client certificate installed in its Trusted Root Certification Authorities Local Computer store.

> **Note** In addition to performing normal certificate validation, the NPS server verifies that the identity sent in the initial EAP-Response/Identity message is the same as the name in the Subject Alternative Name property of the received certificate. This prevents a malicious user from masquerading as a different user or computer from that specified in the EAP-Response/Identity message.

For additional requirements for the wireless client's certificate, see "Requirements for PKI" earlier in this chapter.

By default, NPS performs certificate revocation checking on the certificate received from the wireless clients. You can use the following registry values in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13 on the NPS server to modify certificate revocation checking behavior:

- **IgnoreNoRevocationCheck** When set to 1, NPS accepts EAP-TLS authentications, even when it does not perform or cannot complete a revocation check of the client's certificate chain (excluding the root certificate). Typically, revocation checks fail because the certificate does not include CRL information.
- **IgnoreNoRevocationCheck is set to 0 (disabled) by default. NPS rejects an EAP-TLS or PEAP-TLS authentication unless it can complete a revocation check of the client's certificate chain (including the root certificate) and verify that none of the certificates has been revoked.** Set IgnoreNoRevocationCheck to 1 to accept EAP-TLS or PEAP-TLS authentications when the certificate does not include CRL distribution points, such as those from third-party CAs.
- **IgnoreRevocationOffline** When set to 1, NPS accepts EAP-TLS or PEAP-TLS authentications even when a server that stores a CRL is not available on the network. IgnoreRevocationOffline is set to 0 by default. NPS rejects an EAP-TLS or PEAP-TLS authentication unless it can access CRLs and complete a revocation check of their certificate chain and verify that none of the certificates has been revoked. When it cannot connect to a location that stores a CRL, EAP-TLS or PEAP-TLS considers the certificate to have failed the revocation check.

  Set IgnoreRevocationOffline to 1 to prevent certificate validation failure because of poor network conditions that inhibit revocation checks from completing successfully.

- **NoRevocationCheck** When set to 1, NPS does not perform a revocation check on the wireless client's certificate. The revocation check verifies that the wireless client's certificate and the certificates in its certificate chain have not been revoked. NoRevocationCheck is set to 0 by default.
- **NoRootRevocationCheck** When set to 1, NPS does not perform a revocation check of the wireless client's root CA certificate. This entry eliminates only the revocation check of the client's root CA certificate. A revocation check is still performed on the remainder of the wireless client's certificate chain. NoRootRevocationCheck is set to 0 by default.

  You can use NoRootRevocationCheck to authenticate clients when the root CA certificate does not include CRL distribution points, such as those from third-party CAs. Also, this entry can prevent certification-related delays that occur when a certificate revocation list is offline or is expired.

All these registry values must be added as a DWORD type (a registry data type composed of hexadecimal data with a maximum allotted space of 4 bytes) and set to 0 or 1. The Windows wireless client does not use these values.

Validating the NPS Server's Certificate For the wireless client to validate the certificate of the NPS server, the following must be true for each certificate in the certificate chain sent by the NPS server:

- **The current date must be within the validity dates of the certificate.** When certificates are issued, they are issued with a range of valid dates before which they cannot be used and after which they are considered expired.
- **The certificate has a valid digital signature.** CAs digitally sign certificates they issue. The wireless client verifies the digital signature of each certificate in the chain with the exception of the root CA certificate by obtaining the public key from the certificate's issuing CA and mathematically validating the digital signature.

Additionally, the NPS server computer certificate must have the Server Authentication EKU, which has the object identifier (OID) 1.3.6.1.5.5.7.3.1. To view the EKU for a certificate in the Certificates snap-in, in the contents pane, double-click the certificate, click the Details tab, and then click the Enhanced Key Usage field.

Finally, to trust the certificate chain offered by the NPS server, the wireless client must have the root CA certificate of the issuing CA of the NPS server certificate installed in its Trusted Root Certification Authorities Local Computer store.

For additional requirements for the computer certificate of the NPS server, see "Requirements for PKI" earlier in this chapter.

Notice that the wireless client does not perform certificate revocation checking for the certificates in the certificate chain of the NPS server's computer certificate. The assumption is that the wireless client does not yet have a connection to the network and therefore cannot access a Web page or other resource in order to check for certificate revocation.

Troubleshooting Password-Based Validation

Troubleshooting password validation with PEAP-MS-CHAP v2 authentication consists of verifying the wireless client's user name and password credentials and the computer certificates of the NPS servers.

Validating the Wireless Client's Credentials When you are using PEAP-MS-CHAP v2 for authentication, the name and password as sent by the wireless client must match the credentials of a valid account. The successful validation of the MS-CHAP v2 credentials by the NPS server depends on the following:

- The domain portion of the name corresponds to a domain that is either the domain of the NPS server or a domain that has a two-way trust with the domain of the NPS server.
- The account portion of the name corresponds to a valid account in the domain.
- The password is the correct password for the account.

To verify user account credentials, have the user of the wireless client log on to his or her domain using a computer that is already connected to the network, such as with an Ethernet connection (if possible). This process demonstrates whether there is a problem with the user's credentials or if the problem lies in the configuration of the authentication infrastructure.

Validating the NPS Server's Certificate For the wireless client to validate the certificate of the NPS server for PEAP-MS-CHAP v2 authentication, the following must be true for each certificate in the certificate chain sent by the NPS server:

- **The current date must be within the validity dates of the certificate.** When certificates are issued, they are issued with a valid date range before which they cannot be used and after which they are considered expired.
- **The certificate has a valid digital signature.** CAs digitally sign certificates they issue. The wireless client verifies the digital signature of each certificate in the chain, with the exception of the root CA certificate, by obtaining the public key from the certificate's issuing CA and mathematically validating the digital signature.

Additionally, the NPS server computer certificate must have the Server Authentication EKU (OID 1.3.6.1.5.5.7.3.1). To view the EKU for a certificate in the Certificates snap-in, in the contents pane, double-click the certificate, and then on the Details tab, click the Enhanced Key Usage field.

Finally, to trust the certificate chain offered by the NPS server, the wireless client must have the root CA certificate of the issuing CA of the NPS server certificate installed in its Trusted Root Certification Authorities Local Computer store.

For additional requirements for the computer certificate of the NPS server, see "Requirements for PKI" earlier in this chapter.

Chapter Summary

Deploying a protected wireless network solution involves configuration of Active Directory, PKI, Group Policy, and RADIUS elements of a Windows-based authentication infrastructure and wireless APs and wireless clients. Once deployed, ongoing maintenance consists of managing wireless APs and their configuration for changes in infrastructure servers and updating and deploying wireless profiles. Common problems with wireless connections include the inability to connect due to an authentication or authorization failure and the inability to reach intranet resources from the wireless client.

Additional Information

For additional information about wireless support in Windows Server 2008 and Windows Vista, see the following:

- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Microsoft Wireless Networking (*http://www.microsoft.com/wifi*)

For additional information about Active Directory, see the following:

- Chapter 9, "Authentication Infrastructure"
- *Windows Server 2008 Active Directory Resource Kit*, available as a stand-alone title or in the *Windows Server 2008 Resource Kit* (both Microsoft Press, 2008)
- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support

For additional information about PKI, see the following:

- Chapter 9, "Authentication Infrastructure"
- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/*2008
- Windows Server 2008 Help and Support
- Public Key Infrastructure for Microsoft Windows Server (*http://www.microsoft.com/pki*)
- *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008)

For additional information about Group Policy, see the following:

- Chapter 9, "Authentication Infrastructure"
- *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* (Microsoft Press, 2008)
- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Microsoft Windows Server Group Policy (*http://www.microsoft.com/gp*)

For additional information about RADIUS and NPS, see the following:

- Chapter 9, "Authentication Infrastructure"
- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Microsoft Network Policy Server (*http://www.microsoft.com/nps*)

For additional information about NAP and 802.1X Enforcement, see the following:

- Chapter 14, "Network Access Protection Overview"
- Chapter 15, "Preparing for Network Access Protection"
- Chapter 17, "802.1X Enforcement"
- Windows Server 2008 Technical Library at *http://technet.microsoft.com/windowsserver/2008*
- Windows Server 2008 Help and Support
- Network Access Protection (*http://www.microsoft.com/nap*)