

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

*Joseph Davies and Tony
Northrup with the Microsoft
Networking Team*

PREVIEW CONTENT This excerpt contains uncorrected manuscript from an upcoming Microsoft Press title, for early preview, and is subject to change prior to release. This excerpt is from *Windows Server® 2008 Networking and Network Access Protection (NAP)* from Microsoft Press (ISBN 978-0-7356-2422-1, copyright 2008 Microsoft Corporation, all rights reserved), and is provided without any express, statutory, or implied warranties

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

Microsoft®
Press

978-0-7356-2422-1

© 2008 Microsoft Corporation. All rights reserved.

Table of Contents

Introduction: Infrastructures for Windows Server Networking

Addressing and Packet Flow Infrastructure

Name Resolution Infrastructure

Network Access Infrastructure

Network Access Protection Infrastructure

Part I Addressing and Packet Flow Infrastructure

1 IPv4

Concepts

IPv4 Addressing

IPv4 Routing

Planning and Design Considerations

IPv4 Addressing Scheme

IPv4 Routing Infrastructure

IPv4 Multicast Support

Deployment Steps

IPv4 Routers

Ongoing Maintenance

Adding and Removing Subnets

Troubleshooting

Tools

Common Problems

2 IPv6

Concepts

IPv6 Addressing

IPv6 Routing

IPv6 Transition Technologies

Planning and Design Considerations

Native IPv6

ISATAP

Deployment Steps

Native IPv6

ISATAP

Ongoing Maintenance

Adding and Removing Subnets

Troubleshooting

Tools

Common Problems

3 Dynamic Host Configuration Protocol

Concepts

- DHCP Clients, Servers, and Relay Agents

- DHCP Messages and Options

Planning and Design Considerations

- DHCP Relay Agents

- DHCP Servers

- DHCP and Network Access Protection

- High Availability for DHCP

- DHCP with DNS and WINS

- DHCP and Active Directory

- Designing Scopes

Deployment Steps

- DHCP Relay Agents

- DHCP Servers

- DHCP Client Configuration

Ongoing Maintenance

- Adding and Removing DHCP Servers

- Migrating DHCP Servers

Troubleshooting

- Tools

- Common Problems

4 Windows Firewall with Advanced Security

Concepts

- Filtering Traffic with Windows Firewall

- Protecting Traffic with IPsec

Planning and Design Considerations

- Exceptions for Windows Firewall

- Protected Traffic with IPsec

Deployment Steps

- Firewall Settings with Group Policy

- Server Isolation

- Domain Isolation

- Secure Server

Ongoing Maintenance

- Adding and Removing Windows Firewall Exceptions

- Adding and Removing IPsec Exemptions

Troubleshooting

- Tools

- Common Problems

5 Policy-Based Quality of Service

Concepts

The Causes of Network Performance Problems

How QoS Can Help

QoS for Outbound Traffic

QoS for Inbound Traffic

QoS Implementation

Planning and Design Considerations

Setting QoS Goals

Planning DSCP Values

Planning Traffic Throttling

Hardware and Software Requirements

Planning GPOs and QoS Policies

QoS Policies for Mobile Users of Computers Running Windows Vista

Deployment Steps

How to Configure QoS by Using Group Policy

How to Configure System-Wide QoS Settings

Ongoing Maintenance

Removing QoS Policies

Editing QoS Policies

Monitoring QoS

Troubleshooting

Analyzing QoS Policies

Verifying DSCP Resilience

Isolating Network Performance Problems

6 Scalable Networking

Concepts

TCP Chimney

Receive-Side Scaling

NetDMA

IPsec offload

Planning and Design Considerations

Designating Servers for Scalable Networking

Deployment Steps

Configuring TCP Chimney

Configuring Receive-Side Scaling

Configuring NetDMA

Configuring IPsec Offload

Ongoing Maintenance

Adding or Removing Scalable Networking Servers

Migrating Scalable Networking Servers

Troubleshooting

Tools
Common Problems

Part II Name Resolution Infrastructure

7 Domain Name System

Concepts

DNS Clients and Servers
DNS Name Resolution

Planning and Design Considerations

DNS Namespace
DNS Servers
DNS Security

Deployment Steps

DNS Servers
DNS Zones
DNS Client Configuration
Securing the DNS Infrastructure
High Availability for DNS Servers

Ongoing Maintenance

Adding and Removing DNS Servers
Maintaining Zones

Troubleshooting

Tools
Common Problems

8 Windows Internet Name Service

Concepts

WINS and NetBIOS Name Resolution

Planning and Design Considerations

WINS Server Placement
WINS with DNS and DHCP
WINS Replication

Deployment Steps

WINS Replication Topology
WINS Client Configuration
Securing WINS Traffic

Ongoing Maintenance

Adding and Removing WINS Servers
Maintaining WINS Replication Topology

Troubleshooting

Tools
Common Problems

Part III Network Access Infrastructure

9 Authentication Infrastructure

- Concepts
 - Active Directory Domain Services
 - Public Key Infrastructure
 - Group Policy
 - RADIUS
 - Planning and Design Considerations
- Active Directory
 - PKI
 - Group Policy
 - RADIUS
 - Deployment Steps
 - Deploying Active Directory
 - Deploying PKI
 - Group Policy
 - RADIUS Servers
 - Using RADIUS Proxies for Cross-Forest Authentication
 - Using RADIUS Proxies to Scale Authentications
 - Ongoing Maintenance
- Active Directory
 - PKI
 - Group Policy
 - RADIUS
 - Troubleshooting Tools
- Active Directory
 - PKI
 - Group Policy
 - RADIUS

10 IEEE 802.11 Wireless Networks

- Concepts
 - Support for IEEE 802.11 Standards
 - Wireless Security
 - Components of 802.11 Wireless Networks
 - Planning and Design Considerations
 - Wireless Security Technologies
 - Wireless Authentication Modes
 - Intranet Infrastructure
 - Wireless AP Placement
 - Authentication Infrastructure
 - Wireless Clients
- PKI

- 802.1X Enforcement with NAP
- Deploying Protected Wireless Access
- Deploying Certificates
- Configuring Active Directory for Accounts and Groups
- Configuring NPS Servers
- Deploying Wireless APs
- Configuring Wireless Clients
- Ongoing Maintenance
- Managing User and Computer Accounts
- Managing Wireless APs
- Updating Wireless XML Profiles
- Troubleshooting
- Wireless Troubleshooting Tools in Windows
- Troubleshooting the Windows Wireless Client
- Troubleshooting the Wireless AP
- Troubleshooting the Authentication Infrastructure

11 IEEE 802.1X-Authenticated Wired Networks

- Concepts
- Components of Wired Networks With 802.1X Authentication
- Planning and Design Considerations
- Wired Authentication Methods
- Wired Authentication Modes
- Authentication Infrastructure
- Wired Clients
- PKI
- 802.1X Enforcement with NAP
- Deploying 802.1X-Authenticated Wired Access
- Deploying Certificates
- Configuring Active Directory for Accounts and Groups
- Configuring NPS Servers
- Configuring 802.1X-Capable Switches
- Configuring Wired Clients
- Ongoing Maintenance
- Managing User and Computer Accounts
- Managing 802.1X-Capable Switches
- Updating Wired XML Profiles
- Troubleshooting
- Wired Troubleshooting Tools in Windows
- Troubleshooting the Windows Wired Client
- Troubleshooting the 802.1X-Capable Switch
- Troubleshooting the Authentication Infrastructure

12 Remote Access VPN Connections

- Concepts
- Components of Windows Remote Access VPNs
- Planning and Design Considerations
- VPN Protocols
- Authentication Methods
- VPN Servers
- Internet Infrastructure
- Intranet Infrastructure
- Concurrent Intranet and Internet Access for VPN Clients
- Authentication Infrastructure
- VPN Clients
- PKI
- VPN Enforcement with NAP
- Additional Security Considerations
- Strong Link Encryption
- VPN Traffic Packet Filtering on the VPN Server
- Firewall Packet Filtering for VPN Traffic
- Multi-Use VPN Servers
- Blocking Traffic Routed from VPN Clients
- Concurrent Access
- Unused VPN Protocols
- Deploying VPN-Based Remote Access
- Deploying Certificates
- Configuring Internet Infrastructure
- Configuring Active Directory for User Accounts and Groups
- Configuring RADIUS Servers
- Deploying VPN Servers
- Configuring Intranet Network Infrastructure
- Deploying VPN Clients
- Ongoing Maintenance
- Managing User Accounts
- Managing VPN Servers
- Updating CM Profiles
- Troubleshooting
- Troubleshooting Tools
- Troubleshooting Remote Access VPNs

13 Remote Access VPN Connections

- Concepts
- Components of Windows Remote Access VPNs
- Planning and Design Considerations
- VPN Protocols

- Authentication Methods
- VPN Servers
- Internet Infrastructure
- Intranet Infrastructure
- Concurrent Intranet and Internet Access for VPN Clients
- Authentication Infrastructure
- VPN Clients
- PKI
- VPN Enforcement with NAP
- Additional Security Considerations
- Strong Link Encryption
- VPN Traffic Packet Filtering on the VPN Server
- Firewall Packet Filtering for VPN Traffic
- Multi-Use VPN Servers
- Blocking Traffic Routed from VPN Clients
- Concurrent Access
- Unused VPN Protocols
- Deploying VPN-Based Remote Access
- Deploying Certificates
- Configuring Internet Infrastructure
- Configuring Active Directory for User Accounts and Groups
- Configuring RADIUS Servers
- Deploying VPN Servers
- Configuring Intranet Network Infrastructure
- Deploying VPN Clients
- Ongoing Maintenance
- Managing User Accounts
- Managing VPN Servers
- Updating CM Profiles
- Troubleshooting
- Troubleshooting Tools
- Troubleshooting Remote Access VPNs

Part IV Network Access Protection Infrastructure

14 Network Access Protection Overview

- The Need for Network Access Protection
- Malware and Its Impact on Enterprise Computing
- Preventing Malware on Enterprise Networks
- The Role of NAP
- Business Benefits of NAP
- Components of NAP
- System Health Agents and System Health Validators
- Enforcement Clients and Servers

- NPS
- Enforcement Methods
- IPSec Enforcement
- 802.1X Enforcement
- VPN Enforcement
- DHCP Enforcement
- How NAP Works
- How IPSec Enforcement Works
- How 802.1X Enforcement Works
- How VPN Enforcement Works
- How DHCP Enforcement Works

15 Preparing for Network Access Protection

- Evaluation of Your Current Network Infrastructure
- Intranet Computers
- Networking Support Infrastructure
- NAP Health Policy Servers
- Planning and Design Considerations
- Deployment Steps
- Ongoing Maintenance
- Health Requirement Policy Configuration
- Components of a Health Requirement Policy
- How NAP Health Evaluation Works
- Planning and Design Considerations for Health Requirement Policies
- Remediation Server Groups
- Remediation Server Groups and NAP Enforcement Methods
- Planning and Design Considerations for Remediation Server Groups

16 IPSec Enforcement

- Overview of IPSec Enforcement
- IPSec Enforcement Logical Networks
- Communication Initiation Processes with IPSec Enforcement
- Connection Security Rules for IPSec Enforcement
- Planning and Design Considerations
- Active Directory
- PKI
- HRAs
- IPSec Policies
- NAP Clients
- Deploying IPSec Enforcement
- Configuring Active Directory
- Configuring PKI
- Configuring HRAs
- Configuring NAP Health Policy Servers

- Configuring Remediation Servers on the Boundary Network
- Configuring NAP Clients
- IPSec Enforcement Deployment Checkpoint for Reporting Mode
- Configuring and Applying IPSec Policies
- Ongoing Maintenance
- Adding a NAP Client
- Managing NAP CAs
- Managing HRAs
- Troubleshooting
- Troubleshooting Tools
- Troubleshooting IPSec Enforcement

17 802.1X Enforcement

- Overview of 802.1X Enforcement
- Using an ACL
- Using a VLAN
- Planning and Design Considerations
- Security Group for NAP Exceptions
- 802.1X Authentication Methods
- Type of 802.1X Enforcement
- 802.1X Access Points
- NAP Clients
- Deploying 802.1X Enforcement
- Configuring Active Directory
- Configuring a PEAP-Based Authentication Method
- Configuring 802.1X Access Points
- Configuring Remediation Servers on the Restricted Network
- Configuring NAP Health Policy Servers
- Configuring NAP Clients
- 802.1X Enforcement Deployment Checkpoint for Reporting Mode
- Testing Restricted Access
- Configuring the Network Policy for Noncompliant NAP Clients for Deferred Enforcement
- Configuring Network Policy for Enforcement Mode
- Ongoing Maintenance
- Adding a NAP Client
- Managing 802.1X Access Points
- Troubleshooting
- Troubleshooting Tools
- Troubleshooting 802.1X Enforcement

18 VPN Enforcement

- Concepts

- Planning and Design Considerations

 - Health Policy

- Deployment Steps

 - VPN Servers

 - VPN Enforcement Clients

- Ongoing Maintenance

 - Maintaining health policy for VPN Enforcement

- Troubleshooting

 - Tools

 - Common Problems

19 DHCP Enforcement

- Concepts

- Planning and Design Considerations

 - Health policy

- Deployment Steps

 - DHCP Servers

 - DHCP Enforcement Clients

- Ongoing Maintenance

 - Maintaining health policy for DHCP Enforcement

- Troubleshooting

 - Tools

 - Common Problems

Index

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 Windows Server 2008 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 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 accounts with the Active Directory Users and Computers snap-in. Each user account must be unique.

Computers running Windows Vista, Microsoft Windows XP, Windows Server 2008, or Windows Server 2003 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. Each computer account within a domain must be unique.

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 are 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 for the user and computer account 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.




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

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

Network Access Permission

- You can use this property to set network access permission to be explicitly allowed, denied, or determined through NPS network policies. In all cases, 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 members to the group 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) or 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 approach 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 need to specify only the WirelessUsers group name.

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 publish an up-to-date certificate revocation list (CRL). For more information about CRLs, see the “Certificate Revocation” section of 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 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 Vista and Windows Server 2008, if you trust a root CA (by having 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 a very 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 may be issued for a number of purposes, such as securing e-mail and network authentication. The issuing policy for these uses may be distinct, and separation provides a basis for administering these policies.*

Organizational divisions

- *There may 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 may have entities at multiple physical sites. Network connectivity between these sites may 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 between 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, such as key strength, physical protection, and protection against network attacks.*

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 may 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.

The certification path for a certificate can be viewed on the Certification Path tab of the properties of a certificate by using the Certificates snap-in.

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 should deploy at least 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 to compromise its private key by malicious users. The offline root and intermediate CAs are not required to be Windows Server 2008– or Windows Server 2003–based CAs. Issuing CAs can be subordinates of a third-party intermediate CA. Figure 9-2 shows the recommended enterprise network certificate hierarchy.




Figure 9-2 Recommended certificate hierarchy for enterprise networks

Certificate Revocation

Revocation of a certificate invalidates a 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 get 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, including 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 Institute of Electrical and Electronic Engineers (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 may 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 to store computer and user certificates and manage them by using the Certificates snap-in, subject to Windows security and permissions.*
- *Windows Server 2008 and Windows Server 2003 include Certificate Services, which allows 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 the information in the request and the identity of the requester, issue certificates, revoke certificates, and publish CRLs.

Certificate Services can also be used 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 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

- *A standalone CA is less automated for a user 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 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 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 Group Policy Management Console (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 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 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

Administrators use Group Policy to specify specific configurations for groups of users and computers by creating Group Policy settings. These settings are specified through the GPMC or 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 or 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 options for administrators to set 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 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 zero 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.

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 obviously 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 the following URL; <http://www.ietf.org/>

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, using 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 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 either sends back 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 Network Policy Server (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 obtain user account properties containing authorization and connection setting information.

The 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 definition of *RADIUS client* and *RADIUS server* becomes 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 (to which the RADIUS message is ultimately destined) or another RADIUS proxy. Therefore, when referring to

RADIUS clients and servers from a RADIUS proxy perspective, a RADIUS client is the RADIUS entity from which it receives RADIUS request messages, and a RADIUS server is the RADIUS entity to which it 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).

Note 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 the following URL; <http://www.ietf.org/>

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 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 either the credentials are not authentic or 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 Extensible Authentication Protocol (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 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 requests messages to another NPS server that can authenticate the user or computer account that is attempting to connect.

For the NPS server 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 of 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 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) or 802.1X enforcement when deploying NAP.*

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

The following are planning and design considerations for your PKI that 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. Alternately, 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 or 802.1X enforcement in NAP, 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 and at least one secondary location that are accessible by all computers, especially the RADIUS servers. The RADIUS servers will first attempt to validate the certificate 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

Performing CRL checking is enabled by default for security reasons. It is possible to modify the behavior of NPS for certification revocation checking. There are special cases in which you may 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

Support Escalation Engineer

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 *Microsoft Windows Group Policy Guide: 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 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 *Microsoft Windows Group Policy Guide: 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 decide which domain to make the NPS server computer a member of. 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 this 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 access point (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 domain name system (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 <<**Enterprise and Datacenter SKU names**>> 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 Interdomain 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 if 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 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 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 without your needing to enable it on the NPS server and access server.*

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 (also known as a demilitarized zone or DMZ), 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

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

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 non-Microsoft 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 Internet Protocol security (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:*
 - *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.*

Direct from the Source

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

Technical Lead

-
- *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” 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 it) and a user account name within the domain (the information to the right of it). 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.

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 and 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 the Password Authentication Protocol (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 very 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 (type REG_DWORD) 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 may 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 that 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 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 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 or 802.1X enforcement in NAP, 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 Automatic Allocation 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, open Domains, and then click the name of your 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, then Public Key Policies, and then Automatic Certificate Request Settings.
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 templates page, click Computer, and then click Next. Your enterprise CA appears on the list.
8. Click the enterprise CA, and then 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 with 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, 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. Then click OK*
6. Click Finish. Certificates (Local Computer) or Certificate (*computername*) appears on the list of selected snap-ins for the new console. Click OK.
8. In the console tree, open the Certificates\Personal folder.
9. Right-click the Personal folder, 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 VPN 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 SSTP connections.

To import a computer certificate with the Certificates snap-in

1. Open the Certificates (Local Computer)\Personal folder.
2. Right-click the Personal folder, 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 does 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 the the following URL;
<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 an enterprise CA from Windows Server 2008 Enterprise Edition, Windows Server 2008 Datacenter Edition, Windows Server 2003 Enterprise Edition, or Windows Server 2003 Datacenter Edition, as an issuing CA 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, click the User template.
6. On the Action menu, 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.
7. In the Template Display Name field, type the name of the new user certificate template (for example, VPNAccess).

8. Make sure that the Publish Certificate In Active Directory check box is selected.
9. Click the Security tab.
10. In the Group Or User Names list, click Domain Users.
11. In the Permissions For Domain Users list, select the Read, Enroll, and Autoenroll permission check boxes, and then click OK.
12. Open the Certification Authority snap-in.
13. In the console tree, open Certification Authority, then your CA's name, and then Certificate Templates.
14. On the Action menu, point to New, and then click Certificate To Issue.
15. Click the name of the newly created user certificate template (for example, VPNAccess), and then click OK.
16. Open the Group Policy Management snap-in.
17. In the console tree, open Domains, and then click the name of your domain to which your CA belongs.
18. 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.
19. 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,
20. In the details pane, double-click Autoenrollment Settings.
21. Click Enroll Certificates Automatically.
22. Select the Renew Expired Certificates, Update Pending Certificates, and Remove Revoked Certificates check boxes.
23. Select the Update Certificates That Use Certificate Templates check box, and then click OK.

Perform steps 17–23 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 running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP that 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 with 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, open the Certificates\Personal folder.
6. Right-click the Personal folder, 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 VPN 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 with the Certificates snap-in

1. Open the Certificates (Current User)\Personal folder.
2. Right-click the Personal folder, 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 have 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, open Certificates (Local Computer), open 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, open 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, open Certificates (Local Computer), open 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, open 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.
13. 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, open Certificates (Local Computer), open 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 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

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

Technical Lead

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 *Microsoft Windows Group Policy Guide: Windows Server 2008 and Windows Vista*, resources at <http://www.microsoft.com/gp>, or Windows Server 2008 Help and Support.

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" 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 15, "Preparing for Network Access Protection"*

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 NPS 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 network 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, open Certificates (Local Computer or Computer Name), 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 Computer Name).
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, or 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 folder should appear as the import location. Click Next, and then click Finish.

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

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. Click Start, point to Administrative Tools, and then click Network Policy Server.
2. In the console tree, right-click NPS (Local), and then click Register Server In Active Directory.
3. In the Network Policy Server dialog box, click OK.

Alternately, you can do one of the following:

- *Run 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 running 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" 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 Structured Query Language (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 of 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 of 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 Microsoft SQL Server.
5. To configure a SQL data source, click Configure.
6. On 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, configure the UDP port numbers for your RADIUS authentication traffic, and in the Accounting section, configure 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, and then 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.
4. Optionally, 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 (RADIUS Client Is NAP-Capable check box).
5. 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 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 Appropriate Network Policies

To evaluate authorization and connection constraints for incoming connection requests, you must configure network policies. The following procedure for configuring a network policy is different depending on the type of network access.

To create a network policy

1. From the console tree of the Network Policy Server snap-in, expand Policies, right-click Network Policies, and then click New.
2. On the Specify Network Policy Name And Connection Type page of the New Network Policy wizard, type the name of the network policy, specify the type of access server or the vendor specific number, and then click Next.
3. On the Specify Conditions page, add and configure the conditions for this network policy as needed. Click Next.
4. On the Specify Access Permission page, click Access Granted or Access Denied, and select whether the dial-in properties of an account specify the access. Click Next.
5. On the Configure Authentication Methods page, enable or configure the appropriate authentication methods. Click Next.

6. On the Configure Constraints page, configure the constraints for this network policy as needed. Click Next.
7. On the Configure Settings page, configure the settings for this network policy as needed. Click Next.
8. On the Completing New Network Access Policy page, click Finish.

If you manage the network access permission of user and computer accounts on a per-account basis, use network policies that specify a connection type. If you manage the network access permission through the network policy (the recommended method), use network policies that specify a connection type and group.

Subsequent chapters of this book will provide step-by-step procedures with the New Network Policy wizard, customized for a specific type of network access or NAP enforcement method.

If the access servers require vendor-specific attributes (VSAs), you must add the VSAs to the 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. The listed RFC can be viewed at the following URL; <http://www.ietf.org/>

1. 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.

2. 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.
3. 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.
4. 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

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/machinename.domain.com*). 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

Technical Lead

Note You do not need to use a RADIUS proxy if you use PEAP-MS-CHAP v2 and pre-Windows 2000-style user names (for example: microsoft\user1).

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 pre-Windows 2000 domain.

- *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.

This configuration requires creating at least five RADIUS shared secrets, described as follows:

- *Because typical access servers allow the configuration of only a single RADIUS shared secret for both their primary and secondary RADIUS servers, one shared secret is needed for each access server and the primary and secondary NPS RADIUS proxies.*
- *Because we copy the configuration of the primary NPS proxy to the secondary NPS RADIUS proxy, the following additional RADIUS shared secrets are needed:*
 - *Between the primary and secondary NPS RADIUS proxies and the primary NPS RADIUS server in the first forest*
 - *Between the primary and secondary NPS RADIUS proxies and the secondary NPS RADIUS server in the first forest*
 - *Between the primary and secondary NPS RADIUS proxies and the primary NPS RADIUS server in the second forest*
 - *Between the primary and secondary NPS RADIUS proxies and the secondary NPS RADIUS server in the second forest*

Configuring the Certificate Infrastructure

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

Configuring the Active Directory Forests for Accounts and Groups

Follow the instructions in the “Deploying Active Directory” section of “Deployment Steps” 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 the steps described in the following sections of “Configuring the Primary NPS Server” in this chapter on a computer in the first forest:

- *“Obtaining and Installing a Computer Certificate”*

- *“Configuring NPS Server Properties”*
- *“Configuring Appropriate Network 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” section of “Configuring the Primary NPS Server” 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 the “Configuring the Secondary NPS Server” section 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 sections of “Configuring the Primary NPS Server” in this chapter on a computer in the second forest:

- *“Obtaining and Installing a Computer Certificate”*
- *“Configuring NPS Server Properties”*
- *“Configuring Appropriate Network 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” section of “Configuring the Primary NPS Server” 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 the “Configuring the Secondary NPS Server” section 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. From 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” 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, open RADIUS Clients and Servers,
2. Right-click Remote RADIUS Server Groups, and then click Add.
3. In the New Remote RADIUS Server Group dialog box, type the group name for the NPS RADIUS servers in the first forest in Group Name (for example: RADIUS Servers in Forest1). Click Next.
4. Click Add.
5. 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.
6. 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.
7. Click OK to add the server to the list of servers in the group.
8. In the New Remote RADIUS Server Group dialog box, click Add.
9. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the secondary NPS RADIUS server in the first forest.
10. 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.
11. 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, open RADIUS Clients and Servers.
2. Right-click Remote RADIUS Server Groups, and then click Add.
3. In the New Remote RADIUS Server Group dialog box, type the group name for the NPS RADIUS servers in the second forest in Group Name (for example: RADIUS Servers in Forest2).

4. Click Add.
5. 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.
6. 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.
7. Click OK to add the server to the list of servers in the group.
8. In the New Remote RADIUS Server Group dialog box, click Add.
9. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the secondary NPS RADIUS server in the second forest.
10. 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.
11. 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, open 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 Forest1). Click Next.
3. On the Specify Conditions page, click Add.
4. From 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), and then click OK.
6. Click Next.
7. On the Specify Connection Request Forwarding page, click Forward Requests To The Following Remote RADIUS Server Group For Authentication and select the remote RADIUS server group for the NPS RADIUS servers in the first forest (for example: RADIUS Servers in Forest1). Click Next.
8. On the Configure Settings page, click Next.
9. 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, open 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. From 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), and then click OK.
6. Click Next.
7. On the Specify Connection Request Forwarding page, click Forward Requests To The Following Remote RADIUS Server Group For Authentication and select the remote RADIUS server group for the NPS RADIUS servers in the second forest (for example: RADIUS Servers in Forest2). Click Next.
8. On the Configure Settings page, click Next.
9. 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 needed 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 fallback (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.

This configuration requires the creation of the following RADIUS shared secrets:

- *A different shared secret is needed between each access server and the set of primary and secondary NPS RADIUS proxies. Because typical access servers allow the configuration of only a single RADIUS shared secret for both their primary and secondary RADIUS servers, and because we copy the configuration of the primary NPS proxy to the secondary NPS RADIUS proxy, we cannot use different shared secrets between an access server and the primary and secondary NPS RADIUS proxies.*

- *A different shared secret is needed between each NPS RADIUS server and the set of primary and secondary NPS RADIUS proxies. Because we copy the configuration of the primary NPS proxy to the secondary NPS RADIUS proxy, we cannot use different shared secrets between the primary and secondary NPS RADIUS proxies and each NPS RADIUS server.*

Configuring the Certificate Infrastructure

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

Configuring Active Directory for Accounts and Groups

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

Configuring NPS as a RADIUS Server on Multiple Computers

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

- *“Obtaining and Installing a Computer Certificate”*
- *“Configuring NPS Server Properties”*
- *“Configuring Appropriate Network 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” section of “Configuring the Primary NPS Server” in this chapter (instead of the access servers, add the primary and secondary NPS RADIUS proxies as RADIUS clients).

Note You 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 process 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. From 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” section of “Configuring the Primary NPS Server” in this chapter.

3. In the console tree of the Network Policy Server snap-in, open 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, open 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. From 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), and then click OK.
16. Click Next.
17. On the Specify Connection Request Forwarding page, click Forward Requests To The Following Remote RADIUS Server Group For Authentication and select the remote RADIUS server group for all of the NPS RADIUS servers in the domain. Click Next.
18. On the Configure Settings page, click Next.
19. 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 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" 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 *Microsoft Windows Group Policy Guide: 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" section of "Deployment Steps" in this book.

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 and tools 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 Help and Support, or the resources at <http://www.microsoft.com/pki>.

Digital certificate and PKI-specific troubleshooting issues and tools 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 *Microsoft Windows Group Policy Guide*, Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/gp>.

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

RADIUS

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

- *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/System. Viewing the authentication attempts in this log is useful in troubleshooting network policies. When you have multiple network policies configured, you can use the system 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 system event log is the most useful tool for troubleshooting failed NPS authentications.

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 boxes 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 (or later) is available as a free download from the Microsoft Download Center at <http://www.microsoft.com/downloads>. Network Monitor 3.1 includes a RADIUS parser, which you can use to view the attributes of a RADIUS message and troubleshoot connection issues.

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*
- *Windows Server 2008 Help and Support*
- *Windows Server Active Directory* (<http://www.microsoft.com/ad>)

For additional information about PKI, see the following:

- *Windows Server 2008 Help and Support*
- *Public Key Infrastructure for Windows Server* (<http://www.microsoft.com/pki>)
- *Microsoft Windows Server 2003 PKI and Certificate Security by Brian Komar with the Microsoft PKI team* (Microsoft Press, 2004) (<http://www.microsoft.com/mspress/books/6745.aspx>)

For additional information about Group Policy, see the following:

- *Microsoft Windows Group Policy Guide: Windows Server 2008 and Windows Vista*
- *Windows Server 2008 Help and Support*
- *Windows Server Group Policy* (<http://www.microsoft.com/gp>)
- *Microsoft Windows Group Policy Guide by Darren Mar-Elia, Derek Melber, William Stanek, and the Microsoft Group Policy Team* (Microsoft Press, 2005) (<http://www.microsoft.com/mspress/books/8763.aspx>)

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

- *Windows Server 2008 Help and Support*
- *Internet Authentication Service (<http://www.microsoft.com/ias>)*
- *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"*
- *Windows Server 2008 Help and Support*
- *Network Access Protection (<http://www.microsoft.com/nap>)*