LESSON 2

Authentication, Authorization, and Accounting

OBJECTIVE DOMAIN MATRIX

Skills/Concepts	MTA Exam Objective	MTA Exam Objective Number
Starting Security with Authentication	Understand user authentication.	2.1
Comparing Rights and Permissions	Understand permissions.	2.2
Using Auditing to Complete the Security Picture	Understand audit policies.	2.4
Using Encryption to Protect Data	Understand encryption.	2.5

KEY TERMS

access control list (ACL)
accounting
Active Directory
administrative share
asymmetric encryption
auditing
authentication
authorization
biometrics
BitLocker To Go
brute force attack
built-in groups
certificate chain
certificate revocation list (CRL)
computer account
decryption
dictionary attack
digital certificate
digital signature
domain controller
domain user
effective permissions
encryption
explicit permission
group
hash function
inherited permission
IP Security (IPsec)
Kerberos
key
local user account
member server
multifactor authentication
nonrepudiation

NTFS
NTFS permission
NTLM
organizational units (OU)
owner
password
permission
personal identification number (PIN)
public key infrastructure (PKI)
registry
right
Secure Sockets Layer (SSL)
Security Account Manager (SAM)
security token
share permissions
shared folder
single sign-on (SSO)
smart card
symmetric encryption
syslog
user account
virtual private network (VPN)

The CIO for your company approaches you to discuss security. During the conversation, he asks you what measures the company has in place to ensure that users can access only what they need and nothing else. You respond by explaining that you have built the organization's security model using the three As: authentication, authorization, and accounting. Unfortunately, he wants to know more about this model. How would you respond?

Starting Security with Authentication

THE BOTTOM LINE

In the world of information security, AAA (authentication, authorization, and accounting) is a leading model for access control. Here, ***authentication*** is the process of identifying an individual, usually based on a username and password. After a user is authenticated, he or she can access network resources based on his or her authorization. ***Authorization*** is the process of giving individuals access to system objects based on their identity. Finally, ***accounting***, also known as ***auditing***, is the process of keeping track of a user's activity while accessing network resources, including the amount of time spent in the network, the services accessed while there, and the amount of data transferred during each session.

Nonrepudiation prevents one party from denying the actions it has carried out. If you have established proper authentication, authorization, and accounting, appropriate mechanisms of nonrepudiation should be in place, and no user should be able to deny the actions he or she has carried out while in your organization's system.

CERTIFICATION READY
Can you list the different methods for authentication?
2.1

Before users can access a computer or a network resource, they will most likely log in to prove they are who they say they are and to see whether they have the required rights and permissions to access the network resources.

Logging in is the process through which you are recognized by a computer system or network so that you can begin a session. A user can authenticate via one or more of the following methods:

- **By using what he or she knows:** For instance, by supplying a password or personal identification number (PIN)
- **By using what he or she owns or possesses:** For example, by providing a passport, smart card, or ID card
- **By proving what he or she is:** For instance, by supplying biometric factors based on fingerprints, retinal scans, voice input, etc.

When two or more authentication methods are used to authenticate someone, a ***multifactor authentication*** system is said to be in place. Of course, a system that uses two authentication methods (such as smart cards and passwords) can be referred to as a two-factor authentication system.

Authenticating with What You Know

For both individual computers and entire networks, the most common method of authentication is the password. A ***password*** is a secret series of characters that enables a user to access a particular file, computer, or program.

USING PASSWORDS

When seeking access to a file, computer, or network, hackers will first attempt to crack passwords by trying obvious possibilities, including the names and birthdays of a user's spouse or children, key words used by the user, or the user's hobbies. If these efforts don't work, most hackers will next attempt ***brute force attacks***, which consist of trying as many possible combinations of characters as time and money permit. A subset of the brute force attack is the ***dictionary attack***, which attempts all words in one or more dictionaries. Lists of common passwords are also typically tested.

To make a password more secure, you need to choose a word that nobody can guess. Therefore, whatever you choose should be long enough and should be considered a strong or complex password. For more information about creating strong passwords, visit the following websites:

http://www.microsoft.com/protect/fraud/passwords/create.aspx

https://www.microsoft.com/protect/fraud/passwords/checker.aspx?WT.mc_id=Site_Link

Because today's computers are much more powerful than the computers of years past (which are often used to crack passwords), some people recommend passwords that are at least 14 characters long. However, remembering long passwords can be cumbersome for some people, and these individuals may write their passwords on a piece of paper near their desk. In these situations, you should start looking for other forms of authentication, such as smart cards or biometrics.

Users should also change their passwords regularly; that way, if a user's password is revealed to someone else, it won't be long until that password is no longer valid. In addition, changing passwords routinely also shortens the amount of time that an individual has to guess your password, because he or she will have to start the entire cracking process all over again once your password is changed.

Microsoft includes password policy settings within group policies so that you can easily enforce standards such as minimum number of characters, minimum level of password complexity, how often users must change their passwords, how often users can reuse passwords, and so on.

Although passwords are the easiest security method to implement and the most popular authentication method, use of passwords also has significant disadvantages, including the likelihood of passwords being stolen, spoofed, and/or forgotten. For example, a hacker might call a company's IT department for support and pretend to be a legitimate user, eventually convincing the department to reset that user's password to whatever he or she requests.

Given such scenarios, it's essential that you establish a secure process to reset all user passwords. For instance, you could establish a self-service process in which a user's identity is verified by asking questions and comparing the answers to responses that have been stored previously, such as the person's birthday, the name of his or her favorite movie, the name of his or her pet, and so on. However, these can be relatively easily guessed by an attacker, determined though low-effort research, or discovered through social engineering.

Accordingly, when resetting passwords, you must have a method to positively identify the user who is asking for the password change. Also, you should not send new passwords via

email because if a user's existing password is compromised, the hacker will likely be able to access the user's email account and obtain the new password as well. To avoid these problems, you could meet face-to-face with the person who is requesting a password change and ask for identification. Unfortunately, with large networks and networks that include multiple sites, this may not be plausible. You could also call back and leave the password on the person's voicemail where he or she will need to provide a PIN to access it, or you could send the password to the user's manager or administrative assistant. In either case, you should have the user reset the password immediately after he or she logs on.

USING A PERSONAL IDENTIFICATION NUMBER (PIN)

A ***personal identification number (PIN)*** is a secret numeric password shared between a user and a system that can be used to authenticate the user to the system. Because they only consist of digits and are relatively short (usually four digits), PINs are used for relatively low-security scenarios, such as gaining access to a system, or in combination with another method of authentication.

Authenticating with What You Own or Possess

A second category of authentication is based on what you own or possess. The most common examples of this type of authentication involve use of digital certificates, smart cards, and security tokens.

A ***digital certificate*** is an electronic document that contains an identity, such as a user or organization name, along with a corresponding public key. Because a digital certificate is used to prove a person's identity, it can also be used for authentication. You can think of a digital certificate as similar to a driver's license or passport that contains a user's photograph and thumbprint so that there is no doubt who that user is.

A ***smart card*** is a pocket-sized card with embedded integrated circuits consisting of nonvolatile memory storage components and perhaps dedicated security logic. Nonvolatile memory is memory that does not forget its content when power is discontinued. This kind of memory may contain digital certificates to prove the identity of the person who is carrying the card, and it may also contain permissions and access information. Because smart cards can be stolen, some do not have any markings on them; this makes it difficult for a thief to identify what the card can be used to access. In addition, many organizations require users to supply passwords or PINs in combination with their smart cards.

A ***security token*** (or sometimes a hardware token, hard token, authentication token, USB token, cryptographic token, or key fob) is a physical device that an authorized computer services user is given to ease authentication. Hardware tokens are typically small enough to be carried in a pocket and are often designed to attach to a user's keychain. Some of these security tokens include a USB connector, RFID functions, or Bluetooth wireless interface to enable transfer of a generated key number sequence to a client system. Some security tokens may also include additional technology, such as a static password or digital certificate built into the security token, much like a smart card. Other security tokens may automatically generate a second code that users must input in order to be authenticated.

Authenticating with What You Are

Biometrics is an authentication method that identifies and recognizes people based on physical traits, such as fingerprints, face recognition, iris recognition, retinal scans, and voice recognition. Many mobile computers include a finger scanner, and it is relatively easy to install biometric devices on doors and cabinets to ensure that only authorized people enter secure areas.

To use biometric devices (see Figure 2-1), you must have a biometric reader or scanning device, software that converts the scanned information into digital form and compares match points, and a database that stores the biometric data for comparison.

Figure 2-1

Finger scanner

To launch the biometric system, you will need to set up a station where an administrator enrolls each user; this includes scanning the biometric feature you want to use for authentication. When selecting a biometric method, you should consider its performance, difficulty, reliability, acceptance, and cost. You also need to look at the following characteristics:

- **False reject rate (false negative):** This is the percentage of authorized users who are incorrectly denied access.
- **False accept rate (false positive):** This is the percentage of unauthorized users who are incorrectly granted access.

Introducing RADIUS and TACACS+

When you buy a new computer and create a local user account and login, you are being authenticated with the username and password. For corporations, computers can be part of the domain, and authentication can be provided by the domain controllers. In other situations, you may need to provide centralized authentication, authorization, and accounting when users need to connect to a network service. Two commonly used protocols that provide these functions are Remote Authentication Dial In User Service (RADIUS) and Terminal Access Controller Access-Control System Plus (TACACS+).

A RADIUS or TACACS+ server resides on a remote system and responds to queries from clients such as VPN clients, wireless access points, routers, and switches. The server then authenticates username/password combinations (authentication), determines whether users are allowed to connect to the client (authorization), and logs the connection (accounting).

RADIUS is a mechanism that allows authentication of dial-in and other network connections, including modem dial-up, wireless access points, VPNs, and web servers. As an IETF standard, it has been implemented by most major operating system manufacturers, including Microsoft. For example, in Windows Server 2008, Network Policy Server (NPS) can be used as a RADIUS server to perform authentication, authorization, and accounting for RADIUS clients. It can be configured to use a Microsoft Windows NT Server 4.0 domain, an Active Directory Domain Services (AD DS) domain, or the local Security Accounts Manager (SAM) user accounts database to authenticate user credentials for connection attempts. NPS uses the dial-in properties of the user account and network policies to authorize a connection.

Another competing centralized AAA server is TACACS+, which was developed by Cisco. When designing TACACS+, Cisco incorporated much of the existing functionality of

RADIUS and extended it to meet their needs. From a features viewpoint, TACACS+ can be considered an extension of RADIUS.

Using Run As

Because administrators have full access to individual computers or entire networks, it is recommended that you use a standard nonadministrator user account to perform most tasks. Then, when you need to perform administrative tasks, you can use the Run as command or the built-in options that are included with the Windows operating system.

In previous versions of Windows, you had to use an administrator account to do certain things, such as changing system settings or installing software. If you were logged on as a limited user, the Run as command eliminated the need to log off and then log back on as an administrator.

In newer versions of Windows, including Windows 7 and Windows Server 2008 R2, the Run as command has been changed to Run as administrator. With User Account Control (UAC), you will rarely have to use the Run as administrator command, because Windows automatically prompts you for an administrator password when needed. UAC is discussed in detail in Lesson 5.

RUN A PROGRAM AS AN ADMINISTRATOR

GET READY. To run a program as an administrator, perform the following steps:

1. Right-click the program icon or file that you want to open, and then click **Run as administrator**. See Figure 2-2.
2. Select the administrator account that you want to use, type the password, and then click **Yes**.

You can also use the runas.exe command. For example, to run the widget.exe as an administrator, you would enter the following command:

```
runas /user:admin /widget.exe
```

Figure 2-2

Using the Run as administrator option

■ Introducing Directory Services with Active Directory

THE BOTTOM LINE

A directory service stores, organizes, and provides access to information in a directory. It is used for locating, managing, and administering common items and network resources, such as volumes, folders, files, printers, users, groups, devices, telephone numbers, and other objects. One popular directory service used by many organizations is Microsoft's Active Directory.

Active Directory is a technology created by Microsoft that provides a variety of network services, including the following:

- Lightweight Directory Access Protocol (LDAP)
- Kerberos-based and single sign-on (SSO) authentication
- DNS-based naming and other network information
- A central location for network administration and delegation of authority

The Lightweight Directory Access Protocol, or LDAP, is an application protocol for querying and modifying data using directory services running over TCP/IP. Within the directory, the set of objects is organized in a logical hierarchical manner so that you can easily find and manage those objects. The structure can reflect geographical or organizational boundaries, although it tends to use DNS names for structuring the topmost levels of the hierarchy. Deeper inside the directory, there might be entries representing people, organizational units, printers, documents, groups of people, or anything else that represents a given tree entry (or multiple entries). LDAP uses TCP port 389.

Kerberos is the default computer network authentication protocol, which allows hosts to prove their identity over a nonsecure network in a secure manner. It can also provide mutual authentication so that both the user and server verify each other's identity. To ensure security, Kerberos protocol messages are protected against eavesdropping and replay attacks.

Single sign-on (SSO) allows you to log on once and access multiple related but independent software systems without having to log in again. As you log on with Windows using Active Directory, you are assigned a token, which can then be used to sign on to other systems automatically.

Finally, Active Directory allows you to organize all of your network resources—including users, groups, printers, computers, and other objects—so that you can assign passwords, permissions, rights, and so on to the identity that needs it. You can also assign who is permitted to manage a group of objects.

Looking at Domain Controllers

A ***domain controller*** is a Windows server that stores a replica of the account and security information of a domain and defines the domain boundaries. To make a computer running Windows Server 2008 a domain controller, you will first have to install Active Directory Domain Services. You will then have to execute the dcpromo (short for dc promotion) command to make the server a domain controller from the Search programs and files box, or from the command prompt.

After a computer has been promoted to a domain controller, there are several MMC snap-in consoles to manage Active Directory, including:

- **Active Directory Users and Computers:** Used to manage users, groups, computers, and organizational units.

- **Active Directory Domains and Trusts:** Use to administer domain trusts, domain and forest functional levels, and user principal name (UPN) suffixes.
- **Active Directory Sites and Services:** Used to administer the replication of directory data among all sites in an Active Directory Domain Services (AD DS) forest.
- **Active Directory Administrative Center:** Used to administer and publish information in the directory, including managing users, groups, computers, domains, domain controllers, and organizational units. Active Directory Administrative Center is new in Windows Server 2008 R2.
- **Group Policy Management Console (GPMC):** Provides a single administrative tool for managing Group Policy across the enterprise. GPMC is automatically installed in Windows Server 2008 and newer domain controllers and needs to be downloaded and installed on Windows Server 2003 domain controllers.

Although these tools are typically installed on domain controllers, they can also be installed on client PCs so that you can manage Active Directory without logging on to a domain controller.

Active Directory uses multimaster replication, which means that there is no master domain controller, commonly referred to as a primary domain controller in Windows NT domains. However, there are certain functions that can only be handled by one domain controller at a time.

One role is the PDC Emulator, which provides backwards compatibility for NT4 clients, which is uncommon. However, it also acts as the primary authority for password changes and acts as the master time server within the domain.

A server that is not running as a domain controller is known as a ***member server***. To demote a domain controller to a member server, you would rerun the dcpromo program.

Introducing NTLM

Although Kerberos is the default authentication protocol for today's domain computers, ***NTLM*** is the default authentication protocol for Windows NT, stand-alone computers that are not part of a domain, and situations in which you are authenticating to a server using an IP address. NTLM also acts as a fall-back authentication protocol if Kerberos authentication cannot be completed, such as when it is blocked by a firewall.

NTLM uses a challenge-response mechanism for authentication in which clients are able to prove their identities without sending a password to the server. After a random eight-byte challenge message is sent to the client from the server, the client uses the user's password as a key to generate a response back to the server using an MD4/MD5 hashing algorithm (one-way mathematical calculation) and DES encryption (a commonly used encryption algorithm that encrypted and decrypted data with the same key).

Introducing Kerberos

With Kerberos, security and authentication are based on secret key technology, and every host on the network has its own secret key. The Key Distribution Center maintains a database of these secret keys.

When a user logs in to a network resource using Kerberos, the client transmits the username to the authentication server, along with the identity of the service the user wants to connect to (e.g., a file server). The authentication server constructs a ticket, which randomly generates a key, encrypted with the file server's secret key, and sends it to the client as part of its credentials, which includes the session key encrypted with the client's key. If the user types the right

password, then the client can decrypt the session key, present the ticket to the file server, and give the user the shared secret session key to communicate between them. Tickets are time stamped and typically have an expiration time of only a few hours.

For all of this to work and to ensure security, the domain controllers and clients must have the same time. Windows operating systems include the Time Service tool (W32Time service). Kerberos authentication will work if the time interval between the relevant computers is within the maximum enabled time skew. The default is five minutes. You can also turn off the Time Service tool and install a third-party time service. Of course, if you have problems authenticating, you should make sure that the time is correct for the domain controllers and the client having the problem.

Using Organizational Units

As mentioned earlier, an organization could have thousands of users and thousands of computers. With Windows NT, the domain could only handle so many objects before some performance issues arose. With later versions of Windows, however, the size of the domain was dramatically increased. Whereas with Windows NT you may have required several domains to define your organization, you can now have just one domain to represent a large organization. However, if you have thousands of such objects, you still need a way to organize and manage them.

To help organize objects within a domain and minimize the number of domains, you can use ***organizational units***, or OUs, which can be used to hold users, groups, computers, and other organizational units. See Figure 2-3. An organizational unit can only contain objects that are located in a domain. Although there is no restriction as to how many nested OUs (an OU inside of another OU) you can have, you should design a shallow hierarchy for better performance.

Figure 2-3

Active Directory organizational unit

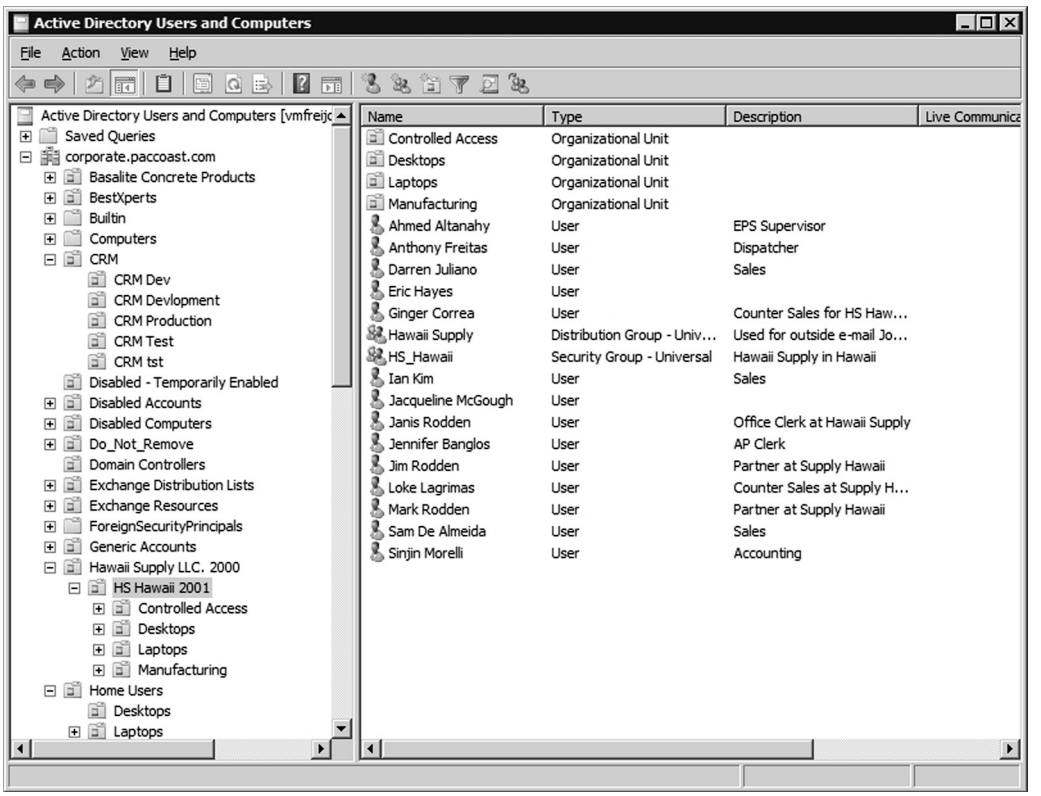

When you first install Active Directory, there are several organizational units already created. They include computers, users, domain controllers, and built-in OUs. Unlike OUs that you create, these OUs do not allow you to delegate permissions or assign group policies. (Group policies will be explained later in the text.) Containers are objects that can store or hold other objects. They include the forest, tree, domain, and organizational unit. To help you manage your objects, you can delegate authority to a container, particularly the domain or organizational unit.

For example, let's say that you have your domain divided by physical location. You can then assign a site administrator authoritative control to the OU that represents a particular physical location, and the user will only have administrative control to the objects within that OU. You can also structure your OUs by function or areas of management. For instance, you could create a Sales OU to hold all of your sales users. You could also create a Printers OU to hold all of the printer objects and then assign a printer administrator to that OU.

Similar to NTFS and the registry, you can assign permissions to users and groups over an Active Directory object. However, you would normally delegate control to the user or group. You can assign basic administrative tasks to regular users or groups and leave domain-wide and forest-wide administration to members of the Domain Admins and Enterprise Admins groups. By delegating administration, you allow groups within your organization to take more control of their local network resources. You also help secure your network from accidental or malicious damage by limiting the membership of administrator groups.

You can delegate administrative control to any level of a domain tree by creating organizational units within a domain, then delegating administrative control for specific organizational units to particular users or groups.

DELEGATE CONTROL

GET READY. To delegate control of an organizational unit, perform the following steps:

1. Open **Active Directory Users and Computers**.
2. In the console tree, right-click the organizational unit for which you want to delegate control.
3. Click **Delegate control** to start the Delegation of Control Wizard, and then follow the instructions.

Looking at Objects

An object is a distinct, named set of attributes or characteristics that represent a network resource. Common objects used within Active Directory are computers, users, groups, and printers. Attributes have values that define the specific object. For example, a user could have the first name John, the last name Smith, and the login name jsmith, all of which identify the user.

When working with objects, administrators typically use the names of those objects, such as usernames. However, all Active Directory objects are also assigned a 128-bit unique number called a security identifier (SID), sometimes referred to as a globally unique identifier (GUID), to uniquely identify them. Therefore, if a user changes his or her username, you can change that name on the network, but he or she will still be able to access all of the same objects and have all of the same rights as before because those objects and rights are assigned to the GUID.

GUIDs also provide some security if a user is deleted. You cannot create a new user account with the same username and expect to have access to all of the objects and all of the rights that the previous user had. Rather, if you decide to let someone in your organization go and

you later replace that person, you should instead disable the first person's account, hire the new person, rename the user account, change the password, and re-enable the account. That way, the new person will be able to access all of the same resources and have all of the same rights that the previous user had.

The schema of Active Directory defines the format of each object and the attributes or fields within each object. The default schema contains definitions of commonly used objects like user accounts, computers, printers, and groups. For example, the schema defines that the user account has fields for first name, last name, and telephone numbers.

To allow Active Directory to be flexible so that it can support other applications, you can extend a schema to include additional attributes. For example, you could add badge number or employee identification fields to the user object. When you install some applications, such as Microsoft Exchange, they will extend the schema, usually by adding additional attributes or fields so that it can support the application.

EXAMINING USERS

A ***user account*** enables a user to log on to a computer and domain. As a result, it can be used to prove the identity of a user, which can then be used to determine what a user can access and what kind of access the user will have (authorization). User accounts can also be used for auditing. For instance, if there is a security problem in which something was inappropriately accessed or deleted, user account data can be used to show who accessed or deleted the object.

On today's Windows networks, there are two types of user accounts:

- Local user account
- Domain user account

A user account allows users to log on and gain access to the computer where the account was created. The ***local user account*** is stored in the ***Security Account Manager (SAM)*** database on the local computer. The only Windows computer that does not have a SAM database is the domain controller. The administrator local user account is the only account that is created and enabled by default in Windows. Although this account cannot be deleted, it can be renamed.

The only other account created by default is the guest account. It was designed for the occasional user who needs access to network resources on a low-security network. The guest local user account is disabled by default and not recommended for general use.

A ***domain user*** account is stored on the domain controller and allows you to gain access to resources within the domain, assuming you have been granted permissions to access those objects. The administrator domain user account is the only account that is created and enabled by default in Windows when you first create a domain. Again, although this account cannot be deleted, it can be renamed.

When you create a domain user account, you must supply a first name, last name, and a user login name. The user login name must be unique with the domain. See Figure 2-4. After the user account is created, you can then open the user account properties and configure a person's username, logon hours, telephone numbers and addresses, which computers the user can log on to, what groups the person is a member of, and so on. You can also specify whether a

Figure 2-4

User account in Active Directory

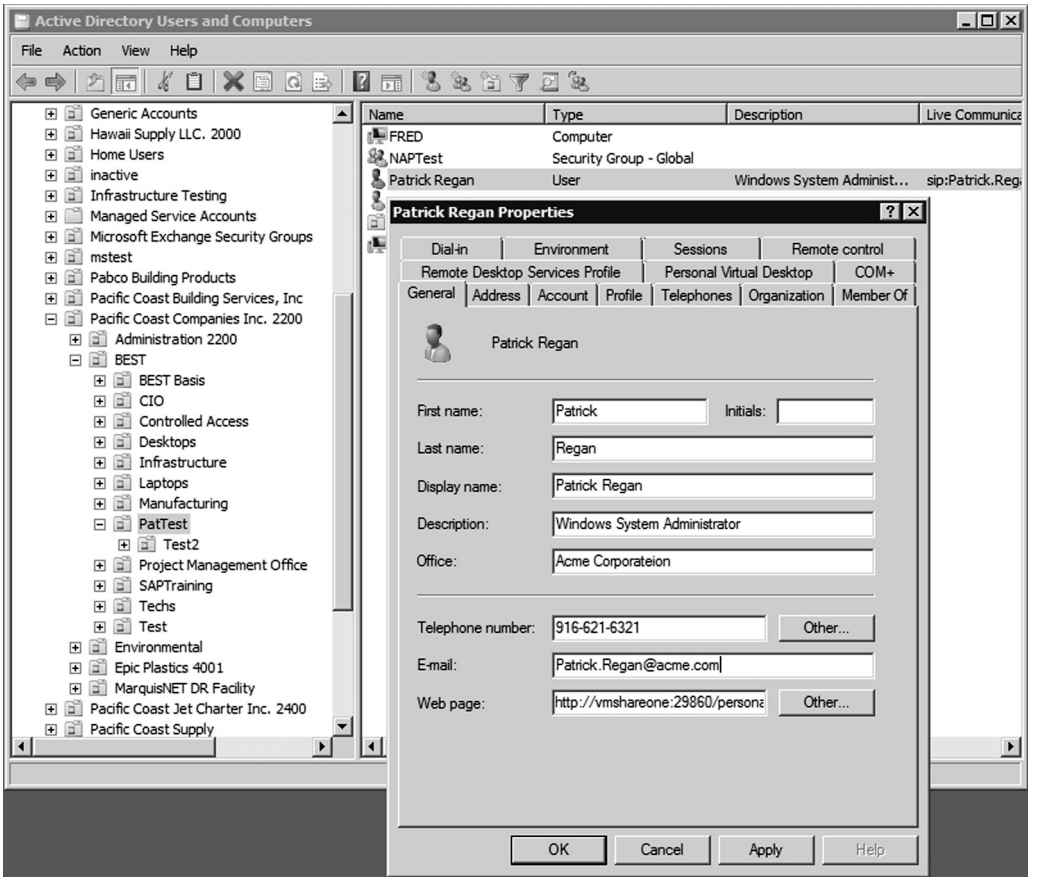

password expires, whether a password can be changed, and whether an account is disabled. Finally, on the Profile tab, you can define the user's home directory, logon script, and profile path. See Figure 2-5.

Figure 2-5

Profile tab

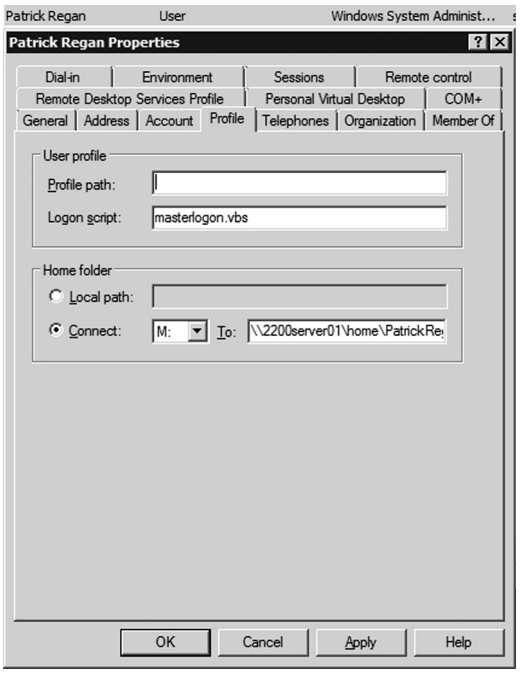

LOOKING AT COMPUTERS

Like user accounts, Windows ***computer accounts*** provide a means for authenticating and auditing a computer's access to a Windows network, as well as its access to domain resources. Each Windows computer to which you want to grant resource access must have a unique computer account. These accounts can also be used for auditing purposes, because they specify what systems were used to access particular resources. See Figure 2-6.

Figure 2-6

Computer account

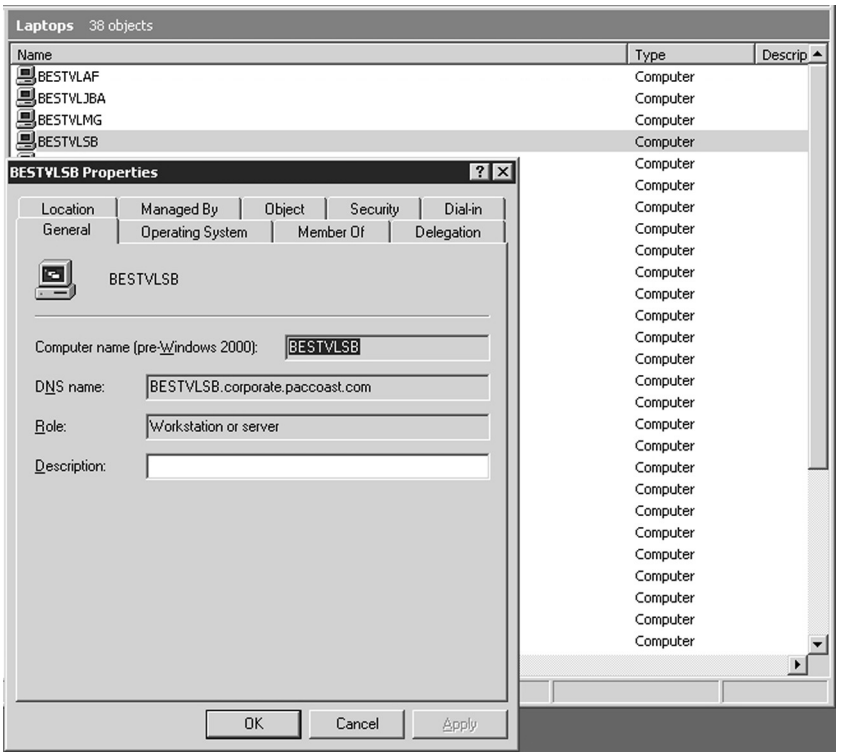

Using Groups

A ***group*** is a collection or list of user accounts or computer accounts. Different from a container, a group does not store users or computers; rather, it just lists them. Using groups can simplify administration, especially when assigning rights and permissions.

A group is used to group users and computers together so that when you assign rights and permissions, you assign them to the group rather than to each user individually. Users and computers can be members of multiple groups, and in some instances, one group can be designated as part of another group.

EXAMINING GROUP TYPES

In Windows Active Directory, there are two types of groups: security and distribution. A security group is used to assign rights and permissions and to gain access to network resources. It can also be used as a distribution group. A distribution group is employed only for nonsecurity functions, such as distributing email, and it cannot be used to assign rights and permissions.

EXAMINING GROUP SCOPES

Any group, whether it is a security group or a distribution group, is characterized by a scope that identifies the extent to which the group is applied in the domain tree or forest. The three group scopes are as follows:

- **Domain local:** Contains global and universal groups, even though it can also contain user accounts and other domain local groups. A domain local group is usually in the domain with the resource to which you want to assign permissions or rights.
- **Global:** Designed to contain user accounts, although they can also contain other global groups. Global groups are designed to be "global" for a domain. After you place user accounts into global groups, these groups are typically placed into domain local groups or universal groups.
- **Universal:** Designed to contain global groups from multiple domains, although they can also contain other universal groups and user accounts. Because global catalogs replicate universal group membership, you should limit membership to global groups. This way, if you change a member within a global group, the global catalog will not have to replicate the change.

See Table 2-1.

Table 2-1

Group scopes

Scope	Members Can Include...	Member Permissions Can Be Assigned...	Group Scope Can Be Converted to...
Universal	Accounts from any domain within the forest in which this universal group resides Global groups from any domain within the forest in which this universal group resides Universal groups from any domain within the forest in which this universal group resides	In any domain or forest	Domain local Global (as long as no other universal groups exist as members)
Global	Accounts from the same domain as the parent global group Global groups from the same domain as the parent global group	In any domain	Universal (as long as the group is not a member of any other global groups)
Domain local	Accounts from any domain, global groups from any domain, universal groups from any domain, and domain local groups but only from the same domain as the parent domain local group	Only within the same domain as the parent domain local group	Universal (as long as no other domain local groups exist as members)

When assigning rights and permissions, you should always try to place your users into groups and assign the rights and permissions to these groups instead of to individual users. To effectively manage the use of global and domain local groups when assigning access to network resources, remember the mnemonic AGDLP (accounts, global, domain local, permissions):

- First, add the user account (A) into the global group (G) in the domain where the user exists.
- Next, add the global group (G) from the user domain into the domain local group (DL) in the resource domain.

- Finally, assign permissions (P) on the resource to the domain local group (DL) in its domain.

If you are using universal groups, the mnemonic is expanded to AGUDLP:

- First, add the user account (A) into the global group (G) in the domain where the user exists.
- Then add the global group (G) from the user domain into the universal group (U).
- Next, add the universal group (U) to the domain local group (DL).
- Finally, assign permissions (P) on the resource to the domain local group (DL) in its domain.

USING BUILT-IN GROUPS

Similar to administrator and guest accounts, Windows has default groups called ***built-in groups***. These default groups have been granted the essential rights and permissions to get you started. Some of Windows' built-in groups are as follows:

- **Domain Admins:** Members of this group can perform administrative tasks on any computer within the domain. By default, the Administrator account is a member.
- **Domain Users:** Windows automatically adds each new domain user account to the Domain Users group.
- **Account Operators:** Members of this group can create, delete, and modify user accounts and groups.
- **Backup Operators:** Members of this group can back up and restore all domain controllers using Windows Backup.
- **Authenticated Users:** This group includes all users with a valid user account on the computer or in Active Directory. Use the Authenticated Users group instead of the Everyone group to prevent anonymous access to a resource.
- **Everyone:** This group includes all users who access a computer with a valid user account.

For more information on the available groups, visit the following website:

http://technet.microsoft.com/en-us/library/cc756898(WS.10).aspx

Looking at Web Server Authentication

When a person accesses a web server, such as those running on Microsoft's Internet Information Server (IIS), several methods of authentication can be used.

When authenticating to web servers, IIS provides a variety of authentication schemes:

- **Anonymous (enabled by default):** Anonymous authentication gives users access to a website without prompting them for a username or password. Instead, IIS uses a special Windows user account called IUSR_*machinename* for access. By default, IIS controls the password for this account.
- **Basic:** Basic authentication prompts the user for a username and password. However, even though the username and password are sent as Base64 encoding, it is basically sent in plain text since Base64 encoding is used as a format and not an encryption. If you need to encrypt usernames and passwords while using basic authentication, you can use digital certificates so that this information is encrypted with https.

- **Digest:** Digest authentication is a challenge/response mechanism that sends a digest or hash using the password as the key instead of sending the password over the network.
- **Integrated Windows authentication:** Integrated Windows authentication (formerly known as NTLM authentication and Windows NT Challenge/Response authentication) can use either NTLM or Kerberos V5 authentication.
- **Client Certificate Mapping:** Client Certificate Mapping uses a digital certificate that contains information about an entity and the entity's public key for authentication purposes.

Comparing Rights and Permissions

THE BOTTOM LINE What a user can do on a system or to a resource is determined by two things: rights and permissions.

CERTIFICATION READY
Can you describe how the permissions are stored for an object?
2.2

A ***right*** authorizes a user to perform certain actions on a computer, such as logging on to a system interactively or backing up files and directories on a system. User rights are assigned through local policies or Active Directory group policies. See Figure 2-7.

Figure 2-7
Group policy user rights assignment

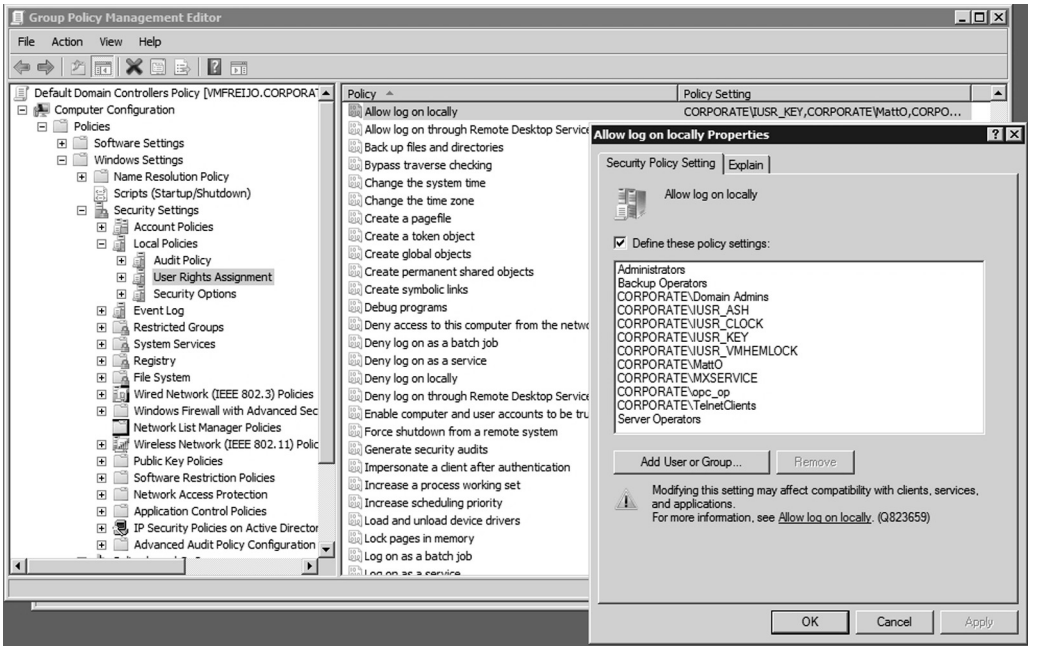

A ***permission*** defines the type of access that is granted to an object (an object can be identified with a security identifier) or object attribute. The most common objects assigned permissions are NTFS files and folders, printers, and Active Directory objects. Information about which users can access an object and what they can do is stored in the ***access control list (ACL)***, which lists all users and groups that have access to an object. NTFS and printer permissions are discussed in the next lesson.

Looking at NTFS

THE BOTTOM LINE

A file system is a method of storing and organizing computer files and the data they contain. It also maintains the physical location of the files so that you can easily find and access the files in the future. Windows Server 2008 supports FAT16, FAT32, and NTFS file systems on hard drives.

After you partition a disk, you then need to format the disk. You can format the disk as FAT16, FAT32, or NTFS. Of these, ***NTFS*** is the preferred file system for today's operating systems.

FAT16, sometimes referred to generically as File Allocation Table (FAT), is a simple file system that uses minimal memory and has been used with DOS. Originally it supported the 8.3 naming scheme, which allowed up to an eight-character filename and three-character filename extension. Later, it was revised to support longer filenames. Unfortunately, FAT volumes can only support up to 2 GB.

FAT32 was released with the second major release of Windows 95. Although this file system can support larger drives, today's Windows supports volumes up to 32 GB. It also supports long filenames.

Today, NTFS is the preferred file system because it supports both volumes up to 16 exabytes and long filenames. In addition, NTFS is more fault tolerant than previous file systems used in Windows because it is a journaling file system. A journaling file system makes sure that a transaction is written to disk properly before being recognized. Finally, NTFS offers better security through permissions and encryption.

Using NTFS Permissions

NTFS permissions allow you to control which users and groups can gain access to files and folders on an NTFS volume. The advantage with NTFS permissions is that they affect both local users and network users.

Usually, when assigning NTFS permissions, you would assign the following standard permissions:

- **Full Control:** Permission to read, write, modify, and execute the files in a folder; change attributes and permissions; and take ownership of the folder or files within
- **Modify:** Permission to read, write, modify, and execute the files in the folder, as well as to change the attributes of the folder or files within
- **Read and Execute:** Permission to display a folder's contents; to display the data, attributes, owner, and permissions for files within the folder; and to run files within the folder
- **List Folder Contents:** Permission to display a folder's contents; and display the data, attributes, owner, and permissions for files within the folder
- **Read:** Permission to display a file's data, attributes, owner, and permissions
- **Write:** Permission to write to a file, append to the file, and read or change the file's attributes

To manage NTFS permissions, you can right-click a drive, folder, or file and select Properties, then select the Security tab. As shown in Figure 2-8, you should see the group and users who have been given NTFS permissions and their respective standard NTFS permissions. To change the permissions, you would click the Edit button.

Figure 2-8
NTFS permissions

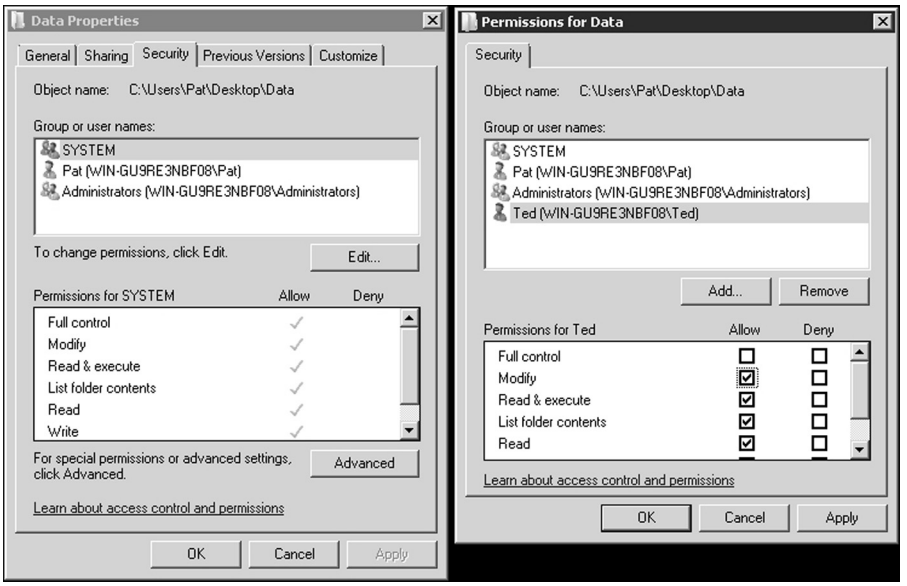

Groups or users who are granted Full Control permission on a folder can delete any files in that folder regardless of the permissions protecting the file. In addition, List Folder Contents is inherited by folders but not files, and it should only appear when you view folder permissions. In Windows Server 2008, the Everyone group does not include the Anonymous Logon group by default, so permissions applied to the Everyone group do not affect the Anonymous Logon group.

To simplify administration, it is recommended that you grant permissions using groups. By assigning NTFS permissions to a group, you are granting permissions to one or more people, reducing the number of entries in each access list and reducing the amount of effort to configure situations in which multiple people need access to certain files or folders.

Looking at Effective NTFS Permissions

The folder/file structure on an NTFS drive can be very complicated and include many folders and many nested folders. In addition, because it is recommended that you assign permissions to groups and at different levels on an NTFS volume, figuring out the effective permissions of a particular folder or file for a particular user can be tricky.

There are two types of permissions used in NTFS:

- ***Explicit permissions:*** Permissions granted directly to a file or folder
- ***Inherited permissions:*** Permissions granted to a folder (parent object or container) that flows into child objects (subfolders or files) inside that folder

When permissions are assigned to a folder, by default, they apply to both the folder and any subfolders and files of that folder. To stop permissions from being inherited in this way, you can select the "Replace all existing inheritable permissions on all descendants with inheritable permissions from this object" in the Advanced Security Settings dialog box. The dialog box will then ask whether you are sure you want to do this. You can also clear the "Allow inheritable permissions from parent to propagate to this object" check box. When the check box is clear, Windows will respond with a Security dialog box. When you click on the Copy button, the explicit permission will be copied from the parent folder to the subfolder or file. You can then change the subfolder's or file's explicit permissions. If you click the Remove button, it will remove the inherited permission altogether.

By default, all objects within a folder inherit the permissions from that folder when they are created. However, explicit permissions take precedence over inherited permissions. So, if you grant different permissions at a lower level, the lower-level permissions will take precedence.

For example, say you have a folder called Data. Within the Data folder, you have Folder 1, and within Folder 1, you have Folder 2. If you grant Allow Full Control to a user account, the Allow Full Control permission will flow down to all the subfolders and files within the Data folder.

Object	NTFS Permissions
Data	Grant Allow Full Control (Explicit)
Folder 1	Allow Full Control (Inherited)
Folder 2	Allow Full Control (Inherited)
File 1	Allow Full Control (Inherited)

Thus, if you grant Allow Full Control on the Data folder to a user account, the Allow Full Control permission will normally flow down to Folder 1. However, if you grant Allow Read permission to Folder 1 to the same user account, the Allow Read permission will overwrite the inherited permission and also flow downward to Folder 2 and File 1.

Object	NTFS Permissions
Data	Grant Allow Full Control (Explicit)
Folder 1	Allow Read (Explicit)
Folder 2	Allow Read (Inherited)
File 1	Allow Read (Inherited)

If a user has access to a file, that user will still be able to gain access to the file even if he or she does not have access to the folder containing the file. Of course, because the user doesn't have access to the folder, the user cannot navigate or browse through the folder to get to the file. Therefore, the user would have to use the universal naming convention (UNC) or local path to open the file.

When you view permissions for an object, they will be one of the following:

- **Checked:** Here, permissions have been explicitly assigned.
- **Cleared (unchecked):** Here, no permissions are assigned.
- **Shared:** Here, permissions are granted through inheritance from a parent folder.

Besides granting Allow permissions, you can also grant the Deny permission. The Deny permission always overrides the other permissions that have been granted, including situations in which a user or group has been given Full Control. For example, if a group has been granted Read and Write permissions yet one member of the group has been denied the Write permission, that user's effective rights would only include the Read permission.

When you combine applying Deny versus Allow permissions and explicit versus inherited permissions, the hierarchy of precedence is as follows:

1. Explicit Deny
2. Explicit Allow
3. Inherited Deny
4. Inherited Allow

Because users can be members of several groups, it is possible for them to have several sets of explicit permissions to a folder or file. When this occurs, the permissions are combined to form the ***effective permissions***, which are the actual permissions when logging in and accessing a file or folder. They consist of explicit permissions plus any inherited permissions.

When you calculate effective permissions, you must first calculate the explicit and inherited permissions for an individual or group and then combine them. When combining user and group permissions for NTFS security, the effective permission is the cumulative permission. The only exception is that Deny permissions always apply.

For example, say you have a folder called Data. Within the Data folder, you have Folder 1, and within Folder 1, you have Folder 2. Imagine also that User 1 is a member of Group 1 and Group 2. If you assign Allow Write permission to the Data folder to User 1, the Allow Read permission to Folder 1 to Group 1, and the Allow Modify permission to Folder 2 to Group 2, then the user's effective permissions would be as follows:

Object	User 1 NTFS Permissions	Group 1 Permissions	Group 2 Permissions	Effective Permissions
Data	Allow Write permission (Explicit)			Allow Write permission
Folder 1	Allow Write permission (Inherited)	Allow Read permission (Explicit)		Allow Read and Write permission
Folder 2	Allow Write permission (Inherited)	Allow Read permission (Inherited)	Allow Modify permission* (Explicit)	Allow Modify permission*
File 1	Allow Write permission (Inherited)	Allow Read permission (Inherited)	Allow Modify permission* (Inherited)	Allow Modify permission*

*The Modify permission includes the Read and Write permissions.

Now, say you have a folder called Data. Within the Data folder, you have Folder 1 and within Folder 1, you have Folder 2. User 1 is a member of Group 1 and Group 2. You assign the Allow Write permission to the Data folder to User 1, the Allow Read permission to Folder 1 to Group 1, and the Deny Modify permission to Folder 2 to Group 2. Here, the user's effective permission would be shown as follows:

Object	User 1 NTFS Permissions	Group 1 Permissions	Group 2 Permissions	Effective Permissions
Data	Allow Write permission (Explicit)			Allow Write permission
Folder 1	Allow Write permission (Inherited)	Allow Read permission (Explicit)		Allow Read and Write permission
Folder 2	Allow Write permission (Inherited)	Allow Read permission (Inherited)	Deny Modify permission (Explicit)	Deny Modify permission
File 1	Allow Write permission (Inherited)	Allow Read permission (Inherited)	Deny Modify permission (Inherited)	Deny Modify permission

VIEW NTFS EFFECTIVE PERMISSIONS

GET READY. To view the NTFS effective permissions granted to a user for a file or folder, perform the following steps:

1. Right-click the file or folder and select **Properties**.
2. Select the **Security** tab.
3. Click the **Advanced** button.
4. Click the **Effective Permissions** tab.
5. Click the **Select** button and type in the name of the user or group you want to view. Click the **OK** button. See Figure 2-9.

Figure 2-9
NTFS effective permissions

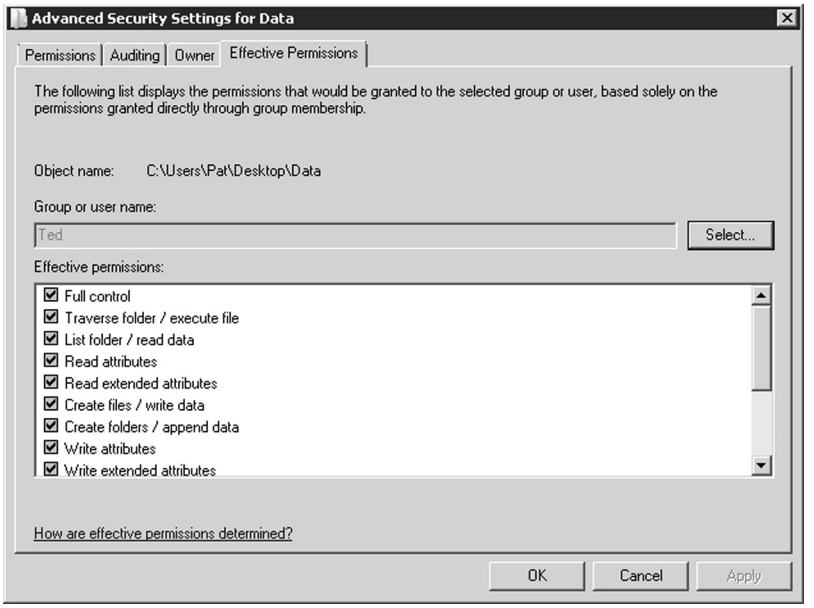

Copying and Moving Files

When you move or copy files from one location to another, you need to understand what happens to the NTFS permissions associated with these files.

When copying and moving files, you will encounter one of three scenarios:

- If you copy a file or folder, the new file or folder will automatically acquire the same permissions as the drive or folder it is being copied to.
- If a file or folder is moved within the same volume, that file or folder will retain the same permissions that were already assigned to it.
- If a file or folder is moved from one volume to another volume, that file or folder will automatically acquire the permissions of the drive or folder it is being copied to.

Using Folder and File Owners

The ***owner*** of an object controls what permissions are set on the object and to whom permissions are granted. If for some reason, you have been denied access to a file or folder and you need to reset the permissions, you can take ownership of the file or folder and then modify the permissions. All administrators automatically have the Take Ownership permission for all NTFS objects.

TAKE OWNERSHIP OF A FILE OR FOLDER

GET READY. To take ownership of a file or folder, perform the following steps:

1. Open **Windows Explorer** and locate the file or folder you want to take ownership of.
2. Right-click the file or folder, click **Properties**, and then click the **Security** tab.
3. Click **Advanced**, and then click the **Owner** tab. See Figure 2-10.

Figure 2-10
Owner tab

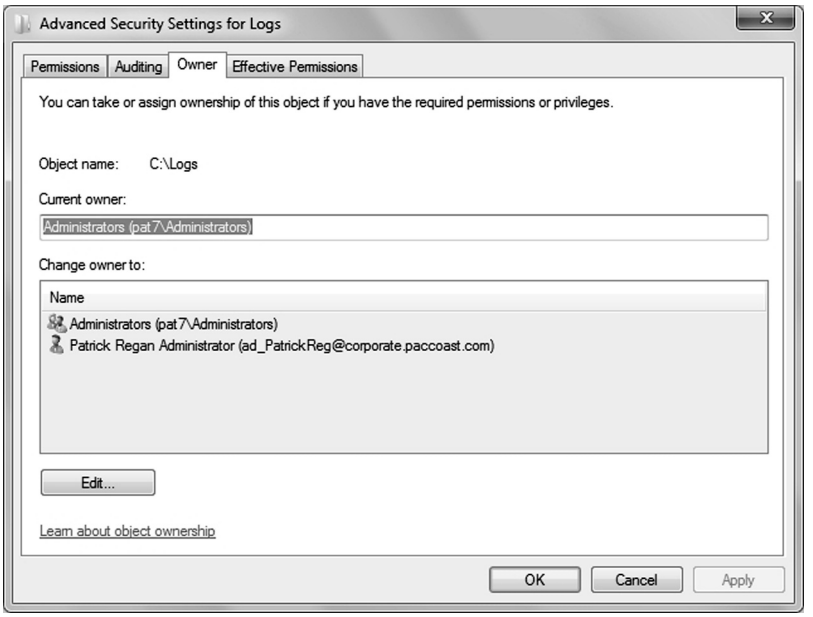

4. Click **Edit**, and then do one of the following:
   - To change the owner to a user or group that is not listed, click **Other users and groups** and, in **Enter the object name to select (examples)**, type the name of the user or group. Then click **OK**.
   - To change the owner to a user or group that is listed, click the name of the new owner in the **Change owner to** box.
5. To change the owner of all subcontainers and objects within the tree, select the **Replace owner on subcontainers and objects** check box.

Sharing Drives and Folders

THE BOTTOM LINE

Most users are not going to log on to a server directly to access their data files. Instead, a drive or folder will be shared (known as a ***shared folder***), and they will access the data files over the network. To help protect against unauthorized access to such folders, you will use share permissions along with NTFS permissions (assuming the shared folder is on an NTFS volume). Then, when users need to access a network share, they will use the Universal Naming Convention UNC, which is \\servername\sharename.

SHARE A FOLDER

GET READY. To share a folder, perform the following steps:

1. In Windows Server 2003, right-click the drive or folder you want to share and select **Sharing and security**. In Windows Server 2008, right-click the drive or folder, select **Properties**, select the **Sharing** tab, and then click the **Advanced Sharing** button.

2. Select **Share this folder**.
3. Type the name of the shared folder.
4. If necessary, you can specify the maximum number of people that can access the shared folder at the same time.
5. Click the **Permissions** button.
6. By default, **Everyone** is given **Allow Read share permission**. Unless you actually want everyone to have access to the folder, you can remove **Everyone**, assign additional permissions, or add additional people.
7. After the desired users and groups have been added with the proper permissions, click the **OK** button to close the **Permissions** dialog box. See Figure 2-11.
8. Click **OK** to close the **Properties** dialog box.

Figure 2-11

Sharing a folder

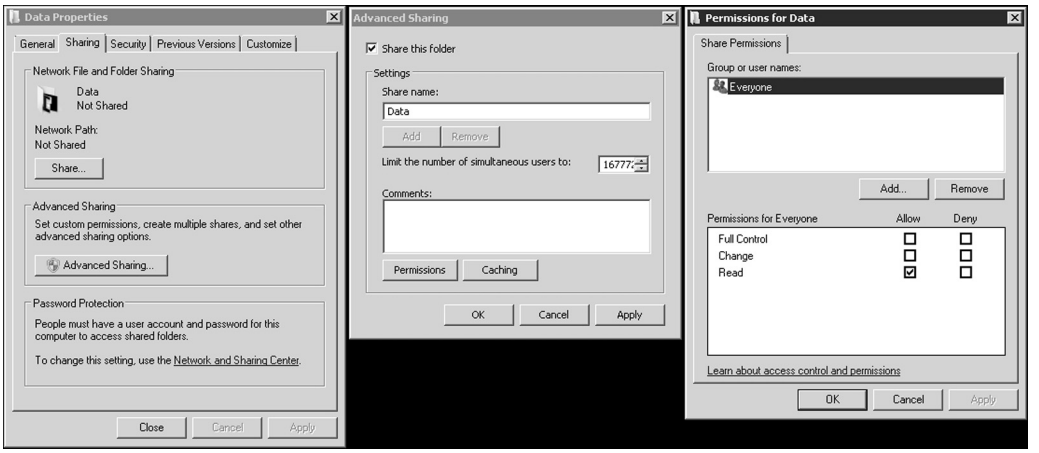

The ***share permissions*** that are available are as follows:

- **Full control:** Users with this permission have Read and Change permissions, as well as the additional capabilities to change file and folder permissions and take ownership of files and folders.
- **Change:** Users with this permission have Read permissions and the additional capabilities to create files and subfolders, modify files, change attributes on files and subfolders, and delete files and subfolders.
- **Read:** Users with this permission can view file and subfolder names, access the subfolders of the share, read file data and attributes, and run program files.

It should be noted that share permissions always apply when accessed remotely using a UNC, even if it is on the FAT, FAT32, or NTFS volume.

As with NTFS, you can also allow or deny each share permission. To simplify managing share and NTFS permissions, Microsoft recommends giving Everyone Full Control, then controlling access using NTFS permissions. In addition, because a user can be member of several groups, it is possible for the user to have several sets of permissions to a shared drive or folder. The effective share permissions are the combination of the user permissions and the permissions for all groups that the user is a member of.

When a person logs directly on to the server console and accesses the files and folders without using the UNC, only the NTFS permissions—and not the share permissions—apply. In contrast, when a person accesses a shared folder using the UNC, you must combine the NTFS and shared permissions to see what a user can do. To determine the overall access, first calculate the effective NTFS permissions, then determine the effective shared permissions. Finally, apply the more restrictive permissions between the NTFS and shared permissions.

Looking at Special and Administrative Shares

In Windows, there are several special shared folders that are automatically created for administrative and system use. Different from regular shares, these shares do not show when a user browses the computer resources using Network Neighborhood, My Network Place, or similar software. In most cases, special shared folders should not be deleted or modified. For Windows Servers, only members of the Administrators, Backup Operators, and Server Operators group can connect to these shares.

An ***administrative share*** is a shared folder typically used for administrative purposes. To make a shared folder or drive into a hidden share, the share name must have a $ at the end of it. Because the share folder or drive cannot be seen during browsing, you would have to use a UNC name that includes the share name (including the $). By default, all volumes with drive letters automatically have administrative shares (C$, D$, E$, and so on). Other administrative shares can be created as needed for individual folders.

Besides the administrative shares for each drive, you will also have the following special shares:

- **ADMIN$:** A resource used by the system during remote administration of a computer. The path of this resource is always the path to the Windows 7 system root (the directory in which Windows 7 is installed—for example, C:\Windows).
- **IPC$:** A resource sharing the named pipes that are essential for communication between programs. It is used during remote administration of a computer and when viewing a computer's shared resources.
- **PRINT$:** A resource used during remote administration of printers.

Introducing the Registry

THE BOTTOM LINE

The ***registry*** is a central, secure database in which Windows stores all hardware configuration information, software configuration information, and system security policies. Components that use the registry include the Windows kernel, device drivers, setup programs, hardware profiles, and user profiles.

Most of the time, you will not need to access the registry because programs and applications typically make all necessary changes automatically. For example, when you change your desktop background or change the default color for Windows, you access the Display settings within the Control Panel and your changes are automatically saved to the registry.

If you do need to access the registry and make changes to it, you should closely follow instructions from a reputable source, because an incorrect change to your computer's registry could render your computer inoperable. However, there may be a time when you need to make a change in the registry because there is no interface or program to make the change. To view and manually change the registry, you will use the Registry Editor (Regedit.exe), which can be executed from the command prompt, Start Search box, or Run box. See Figure 2-12.

Figure 2-12

Registry Editor

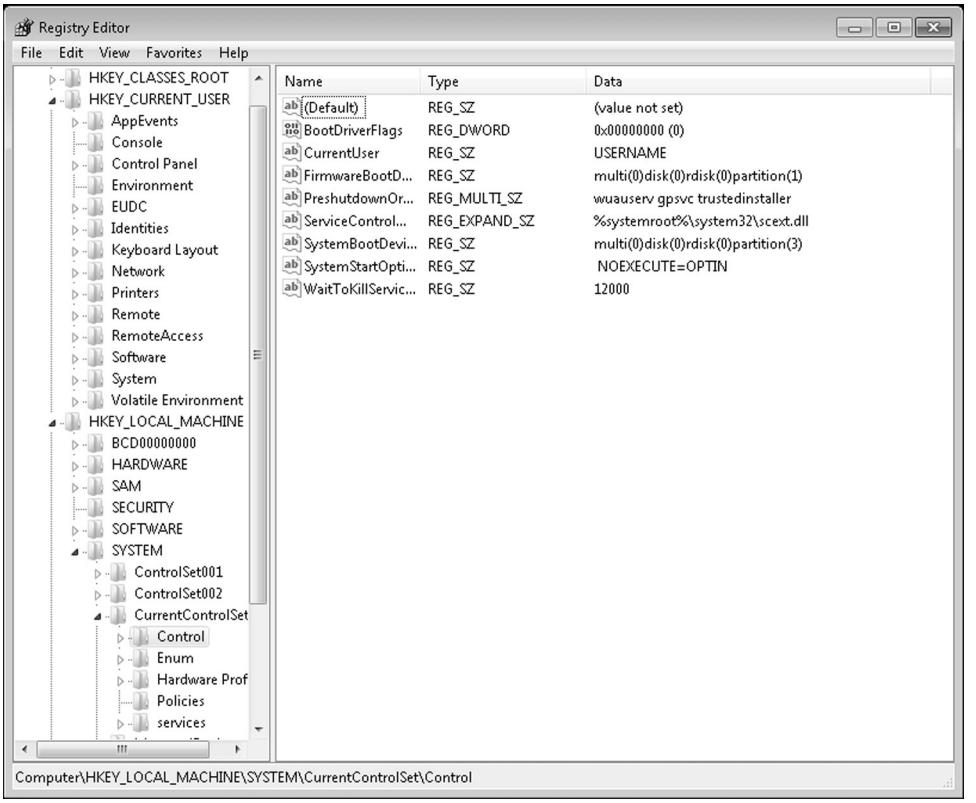

The registry is split into a several logical sections, often referred to as hives, which are generally named by their Windows API definitions. The hives begin with HKEY and are often abbreviated to a three- or four-letter short name starting with "HK." For example, HKCU is HKEY_CURRENT_USER, and HKLM is HKEY_LOCAL_MACHINE. Windows 7 has five Root Keys/HKEYs:

- **HKEY_CLASSES_ROOT:** Stores information about registered applications, such as file association data that tells which default program opens files with a certain extension.
- **HKEY_CURRENT_USER:** Stores settings that are specific to the currently logged-in user. When a user logs off, the HKEY_CURRENT_USER is saved to HKEY_USERS.
- **HKEY_LOCAL_MACHINE:** Stores settings that are specific to the local computer.
- **HKEY_USERS:** Contains subkeys corresponding to the HKEY_CURRENT_USER keys for each user profile actively loaded on the machine.
- **HKEY_CURRENT_CONFIG:** Contains information gathered at run time. Information stored in this key is not permanently stored on disk, but rather regenerated at boot time.

Registry keys are similar to folders that contain values or subkeys. The keys within the registry follow a syntax similar to a Windows folder or file path that uses backslashes to separate each level. For example:

HKEY_LOCAL_MACHINE\Software\Microsoft\Windows

refers to the subkey "Windows" of the subkey "Microsoft" of the subkey "Software" of the HKEY_LOCAL_MACHINE key.

Registry values include a name and a value. There are multiple types of values. Some of the most common key types are shown in Table 2-2.

Table 2-2
Common registry key types

Name	Data Type	Description
Binary value	REG_BINARY	Raw binary data. Most hardware component information is stored as binary data and displayed in Registry Editor in hexadecimal format.
DWORD value	REG_DWORD	Data represented by a number that is four bytes long (a 32-bit integer). Many parameters for device drivers and services are this type and are displayed in Registry Editor in binary, hexadecimal, or decimal format.
Expandable string value	REG_EXPAND_SZ	A variable-length data string. This data type includes variables that are resolved when a program or service uses the data.
Multi-string value	REG_MULTI_SZ	A multiple string. Values that contain lists or multiple values in a form that people can read are generally this type. Entries are separated by spaces, commas, or other marks.
String value	REG_SZ	A fixed-length text string.
QWORD value	REG_QWORD	Data represented by a number that is a 64-bit integer. This data is displayed in Registry Editor as a binary value and was introduced in Windows 2000.

Reg files (also known as registration entries) are text files for storing portions of a registry. These files have a .reg filename extension. If you double-click a reg file, it will add the registry entries into the registry. You can export any registry subkey by right-clicking the subkey and choosing Export. You can back up the entire registry to a reg file by right-clicking Computer at the top of Regedit and selecting export, or you can back up the system state with Windows Backup.

ACCESS REGISTRY PERMISSIONS

GET READY. The registry uses permissions that are stored in Access Control Lists (ACLs). To access the registry permissions, perform the following steps:

1. Open **Registry Editor**.
2. Click the key to which you want to assign permissions.
3. On the **Edit** menu, click **Permissions**.

You will then add the prospective user and assign either Allow or Deny Full Control or Read permission.

■ Using Encryption to Protect Data

THE BOTTOM LINE

Encryption is the process of converting data into a format that cannot be read by another user. Once a user has encrypted a file, that file automatically remains encrypted when it is stored on disk. ***Decryption*** is the process of converting data from an encrypted format back to its original format.

CERTIFICATION READY
Can you list and contrast the three primary methods of encryption?
2.5

With commonly used encryption, the encryption algorithm needs to provide a high level of security yet still be available to the public. Because the algorithm is made available to the public, the security resides in the key, not in the algorithm itself.

One of the simplest cipher algorithms is the substitution cipher, which changes one character or symbol into another. For example, if you have

clear text

and you substitute each "e" with a "y," each "c" with the letter "j," and each letter "t" with a "y," you would get the following cipher text:

jlyar yexy

Another simple technique is based on the transposition cipher, which involves transposing or scrambling letters in a certain manner. For example, if you have

clear text

and you switch every two letters, you get:

lcae rettx

A ***key***, which can be thought of as a password, is applied mathematically to plain text to provide cipher or encrypted text. Different keys produce different encrypted output. With computers, encryption is often based on bits, not characters. For example, if you have the Unicode letters "cl," it could be expressed in the following binary format:

01100011 01101100

If you mathematically add the binary form of 'z'(01111010), which is the key, you get:

```
 01100011      01101100
+01111010     +01111010
 11011101      1110 0110
```

which would appear as strange Unicode characters: ýæ.

Like a password, the longer a key is (usually expressed in bits), the more secure it is. For a hacker to figure out a key, he or she would also have to use a brute force attack, which means the hacker would have to try every combination of bits until he or she figured out the correct key. Although a key could be broken given enough time and processing power, long keys are chosen so that key cracking will take months, maybe even years, to accomplish. Of course, as with passwords, some encryption algorithms change their key frequently. Therefore, a key length of 80 bits is generally considered the minimum for strong security with symmetric encryption algorithms. 128-bit keys are commonly used and are also considered very strong.

Examining Types of Encryption

Encryption algorithms can be divided into three classes: symmetric, asymmetric and hash function.

LOOKING AT SYMMETRIC ENCRYPTION

Symmetric encryption uses a single key to encrypt and decrypt data. Therefore, it is also referred to as secret-key, single-key, shared-key, and private-key encryption. To use symmetric key algorithms, you need to initially exchange the secret key between both sender and receiver.

Symmetric-key ciphers can be divided into block ciphers and stream ciphers. A block cipher takes a block of plain text and a key, and then outputs a block of cipher text of the same size. Two popular block ciphers include the Data Encryption Standard (DES) and the Advanced Encryption Standard (AES), which have been designated cryptography standards by the U.S. government.

The National Bureau of Standards selected the Data Encryption Standard as an official Federal Information Processing Standard (FIPS) for the United States in 1976. It is based on a symmetric-key algorithm that uses a 56-bit key.

Because DES is based on a relatively small 56-bit key size, it was subject to brute force attacks. Therefore, instead of designing a completely new block cipher algorithm, Triple DES (3DES), which uses three independent keys, was developed. DES and the more secure 3DES are still popular and used across a wide range of applications, ranging from ATM encryption, to email privacy, to secure remote access.

Although DES and 3DES remain popular, a more secure encryption method called Advanced Encryption Standard (AES) was announced in 2001 and is currently growing in popularity. This standard comprises three block ciphers—AES-128, AES-192, and AES-256—used on 128-bit blocks with key sizes of 128, 192, and 256 bits, respectively. The AES ciphers have been analyzed extensively and are now used worldwide, including with Wi-Fi Protected Access 2 (WPA2) wireless encryption.

In contrast with block ciphers, stream ciphers create an arbitrarily long stream of key material, which is combined bit-by-bit or character-by-character with the plain text. RC4 is one widely used stream cipher, employed in both Secure Sockets Layer (SSL) and Wired Equivalent Privacy (WEP). Although RC4 is simple and known for its speed, it can be vulnerable if the key stream is not discarded, nonrandom or related keys are used, or a single key stream is used twice.

LOOKING AT ASYMMETRIC ENCRYPTION

Asymmetric encryption, also known as public key cryptography, uses two mathematically related keys for encryption. One key is used to encrypt the data, while the second is used to decrypt it. Unlike symmetric key algorithms, this method does not require a secure initial exchange of one or more secret keys to both sender and receiver. Instead, you can make the public key known to anyone and use the other key to encrypt or decrypt the data. The public key could be sent to someone or could be published within a digital certificate via a Certificate Authority (CA). Secure Sockets Layer (SSL)/Transport Layer Security (TLS) and Pretty Good Privacy (PGP) all use asymmetric keys. Two popular asymmetric encryption protocols are Diffie-Hellman and RSA.

For example, say you want a partner to send you data. To begin the asymmetric encryption process, you send your partner the public key. Your partner will then encrypt the data with the key and send you the encrypted message. You will next use the private key to decrypt the message. If the public key falls into someone else's hands, that person still could not decrypt the message because you need the private key to decrypt a message that has been encrypted with the public key.

LOOKING AT HASH FUNCTION

The last type of encryption is the hash function. Different from the symmetric and asymmetric algorithms, a ***hash function*** is meant as a one-way encryption. That means that after something has been encrypted with this method, it cannot be decrypted. For example, a hash function can be used to encrypt a password that is stored on disk and for digital signatures. Anytime a password is entered, the same hash calculation is performed on the entered password and compared to the hash value of the password stored on disk. If the two match, the user must have typed in the password. This avoids storing passwords in a readable format where a hacker might be able to gain access to them.

Introducing Public Key Infrastructure

Public key infrastructure (PKI) is a system consisting of hardware, software, policies, and procedures that create, manage, distribute, use, store, and revoke digital certificates. Within the PKI, the certificate authority (CA) binds a public key with respective user identities and issues digital certificates containing the public key.

For the PKI system to work, the CA must be trusted. Typically within an organization, you may install a CA on Windows server, specifically on a domain controller, and it would be trusted within your organization. If you require a CA that is trusted outside your organization, you would have to use a trusted third-party CA, such as VeriSign or Entrust. Established commercial CAs charge to issue certificates that will automatically be trusted by most web browsers. See Figure 2-13.

Figure 2-13

Trusted CAs in Internet Explorer

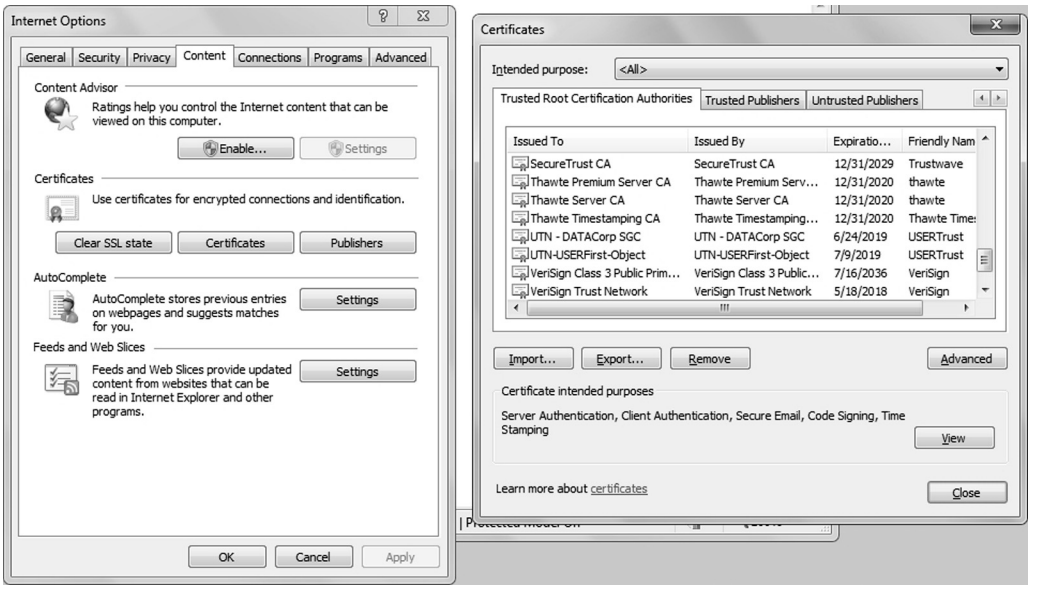

The registration authority (RA), which may or may not be the same server as the CA, is used to distribute keys, accept registrations for the CA, and validate identities. The RA does not distribute digital certificates; instead, the CA does.

Besides having an expiration date, a digital certificate can be revoked if it was compromised or if the situation has changed for the system to which the certificate was assigned. A ***certificate revocation list (CRL)*** is a list of certificates (or more specifically, a list of serial numbers for certificates) that have been revoked or are no longer valid and therefore should not be relied on.

As previously mentioned, Windows servers can host a certificate authority. The Enterprise Root CA is at the top level of the certificate authority hierarchy. Once Enterprise Root CA is configured, it registers automatically within Active Directory, and all computers within the domain trust it. This authority will support auto enrollment and auto-renewal of digital certificates.

If you need to support outside clients and customers, you would most likely build a stand-alone CA. Unlike the Enterprise Root CA, a stand-alone CA does not use Active Directory. Because stand-alone CAs do not support auto enrollment, all requests for certificates are pending until an administrator approves them.

USING DIGITAL CERTIFICATES

A ***digital certificate*** is an electronic document that contains a person's or organization's name, a serial number, an expiration date, a copy of the certificate holder's public key (used for encrypting messages and creating digital signatures), and the digital signature of the CA that assigned the certificate so that recipients can verify that the certificate is real.

The most common digital certificate is the X.509 version 3. The X.509 version 3 standard specifies the format for the public key certificate, certificate revocation lists, attribute certificates, and a certificate path validation algorithm. See Figure 2-14.

Figure 2-14

X.509 digital certificate

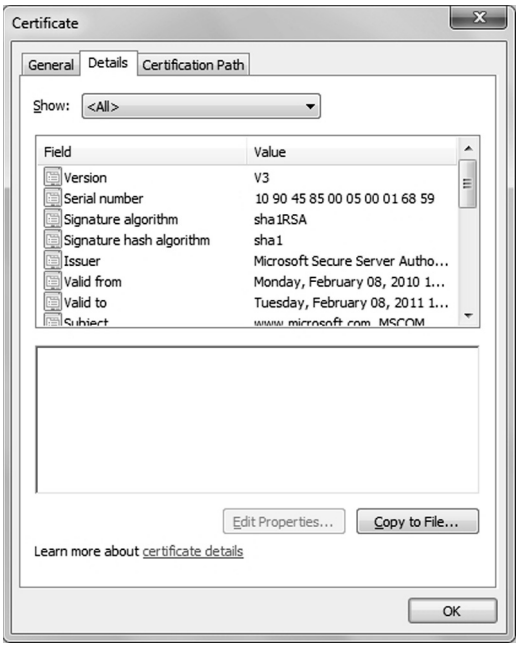

Digital certificates can be imported and exported via electronic files. Four common formats are as follows:

- **Personal Information Exchange (PKCS #12):** The Personal Information Exchange format (PFX, also called PKCS #12) supports secure storage of certificates, private keys, and all certificates in a certification path. The PKCS #12 format is the only file format that can be used to export a certificate and its private key. It will usually have a .p12 filename extension.
- **Cryptographic Message Syntax Standard (PKCS #7):** The PKCS #7 format supports storage of certificates and all certificates in a certification path. It will usually have a .p7b or .p7c filename extension.
- **DER-encoded binary X.509:** The Distinguished Encoding Rules (DER) format supports storage of a single certificate. This format does not support storage of the private key or certification path. It will usually have a .cer, .crt, or .der filename extension.
- **Base64-encoded X.509:** The Base64 format supports storage of a single certificate. This format does not support storage of the private key or certification path.

ACQUIRE A DIGITAL CERTIFICATE

GET READY. To acquire a digital certificate using IIS 7/7.5, perform the following steps:

1. Request an Internet server certificate from the IIS server. To do so, click the server within **IIS Manager**, then double-click **Server Certificates** in the **Features** view. Next click **Create Certificate Request** from the **Actions** pane.

2. Send the generated certificate request to the CA, usually using the vendor's website.
3. Receive a digital certificate from the CA and install it on the IIS server. Again, open **IIS Manager**, double-click the server within IIS Manager, and double-click **Server Certificates** in the **Features** view. Then select **Complete Certificate Request**.

If you have a web farm that consist of multiple web servers, you will need to install the digital certificate from the first server and export the digital certificate to a pfx format, and you will need to copy the public and private key to the other servers. Therefore, you will need to export the key from the first server and import to the other servers.

EXPORT A DIGITAL CERTIFICATE

GET READY. To export a digital certificate, perform the following steps:

1. Open **IIS Manager** and navigate to the level you want to manage.
2. In the **Features** view, double-click **Server Certificates**.
3. In the **Actions** pane, click **Export**.
4. In the **Export** dialog box, type a filename in the **Export to** box or click the **Browse** button to navigate to the name of a file in which to store the certificate for exporting.
5. Type a password in the **Password** box if you want to associate a password with the exported certificate. Retype the password in the **Confirm password** box.
6. Click **OK**.

IMPORT A DIGITAL CERTIFICATE

GET READY. To import a certificate, perform the following steps:

1. Open **IIS Manager** and navigate to the level you want to manage.
2. In the **Features** view, double-click **Server Certificates**.
3. In the **Actions** pane, click **Import**.
4. In the **Import Certificate** dialog box, type a filename in the **Certificate file** box or click the **Browse** button to navigate to the name of the file where the exported certificate is stored. Type a password in the **Password** box if the certificate was exported with a password.
5. Select **Allow this certificate to be exported** if you want to be able to export the certificate, or clear **Allow this certificate to be exported** if you want to prevent additional exports of this certificate.
6. Click **OK**.

EXAMINING A CERTIFICATE CHAIN

There are only so many root CA certificates that are assigned to commercial third-party organizations. Therefore, when you acquire a digital certificate from a third-party organization, you might need to use a certificate chain to obtain the root CA certificate. In addition, you may need to install an intermittent digital certificate that will link the assigned digital certificate to a trusted root CA certificate. The ***certificate chain***, also known as the certification path, is a list of certificates used to authenticate an entity. It begins with the certificate of the entity and ends with the root CA certificate. See Figure 2-15.

Figure 2-15
Certificate chain

USING A DIGITAL SIGNATURE

A ***digital signature*** is a mathematical scheme that is used to demonstrate the authenticity of a digital message or document. It is also used to prove that the message or document has not been modified. With a digital signature, the sender uses the receiver's public key to create a hash of the message, which is stored in the message digest. The message is then sent to the receiver. The receiver will next use his or her private key to decrypt the hash value, perform the same hash function on the message, and compare the two hash values. If the message has not been changed, the hash values will match.

To prove that a message comes from a particular person, you can perform the hashing function with your private key and attach the hash value to the document to be sent. When the document is sent and received by the receiving party, the same hash function is completed. You then use the sender's public key to decrypt the hash value included in the document. If the two hash values match, the user who sent the document must have known the sender's private key, proving who sent the document. It will also prove that the document has not been changed.

USING SECURE SOCKETS LAYER (SSL) AND TRANSPORT LAYER SECURITY (TLS)

There are times when you need to transmit private data over the Internet, such as credit card numbers, Social Security numbers, and so on. In these instances, you should use SSL over http (https) to encrypt the data before sending it. By convention, URLs that require an SSL connection start with https: instead of http:.

SSL is short for ***Secure Sockets Layer***. It is a cryptographic system that uses two keys to encrypt data, a public key known to everyone and a private or secret key known only to the recipient of the message. The public key is published in a digital certificate, which also confirms the identity of the web server.

When you connect to a site that is secured using SSL, a gold lock appears in the address bar, along with the name of the organization to which the CA issued the certificate. Clicking the lock icon displays more information about the site, including the identity of the CA that

issued the certificate. For even more information, you can click the View Certificate link to open the Certificate dialog box.

Occasionally, Internet Explorer may find problems with a website's digital certificate—for instance, the certificate may be expired, may be corrupted, may have been revoked, or may not match the name of the website. When this happens, IE will block access to the site and display a warning stating that there is a problem with the certificate. You then have a chance to close the browser window or ignore the warning and continue on to the site. Of course, if you chose to ignore the warning, make sure you trust the website and you believe that you are communicating with the correct server.

Transport Layer Security (TLS) is an extension of SSL that was supported by the Internet Engineering Task Force (IETF) so that it could be an open, community-supported standard that could then be expanded with other Internet standards. Although TLS is often referred to as SSL 3.0, it does not interoperate with SSL. Also, even though TLS is usually the default for most browsers, it has a downgrade feature that allows SSL 3.0 to run as needed.

Encrypting Email

Because email is sent over the Internet, you may be concerned with the possibility that your data packets will be captured and read. Therefore, you may want to encrypt emails that contain confidential information.

There are multiple protocols that can be used to encrypt emails. Two prominent protocols include:

- Secure Multipurpose Internet Mail Extension (S/MIME)
- Pretty Good Privacy (PGP)

Secure Multipurpose Internet Mail Extension (S/MIME) is the secure version of MIME, used to embed objects within email messages. It is the most widely supported standard used to secure email communications, and it uses the PKCS #7 standard. S/MIME is included with popular web browsers and has also been endorsed by other messaging products vendors.

Pretty Good Privacy (PGP) is a freeware email encryption system that uses symmetrical and asymmetrical encryption. Here, when email is sent, the document is encrypted with the public key and also a session key. The session key is a one-use random number used to create the cipher text. The session key is encrypted into the public key and sent with the cipher text. When the message is received, the private key is used to extract the session key. The session key and the private key are then used to decrypt the cipher text.

Encrypting Files with EFS

If someone steals a hard drive that is protected by NTFS permissions, that person could take the hard drive, put it in a system of which he or she is an administrator, and access all files and folders on the hard drive. Therefore, to truly protect a drive that could be stolen or accessed illegally, you can encrypt the files and folders on that drive.

Windows 7 offers two file encrypting technologies, Encrypting File System (EFS) and BitLocker Drive Encryption. EFS protects individual files or folders, whereas BitLocker protects entire drives.

Encrypting File System (EFS) can encrypt files on an NTFS volume so that they cannot be used unless the user has access to the keys required to decrypt the information. After a file has

been encrypted, you do not have to manually decrypt the encrypted file before you can use it. Rather, once you encrypt a file or folder, you work with the encrypted file or folder just as you would with any other file or folder.

EFS is keyed to a specific user account, using the public and private keys that are the basis of the Windows public key infrastructure (PKI). The user who creates a file is the only person who can read it. As the user works, EFS encrypts the files he or she creates using a key generated from the user's public key. Data encrypted with this key can be decrypted only by the user's personal encryption certificate, which is generated using his or her private key.

ENCRYPT A FOLDER OR FILE USING EFS

GET READY. To encrypt a folder or file, perform the following steps:

1. Right-click the folder or file you want to encrypt, then click **Properties**.
2. Click the **General** tab, and then click **Advanced**.
3. Select the **Encrypt contents to secure data** check box, click **OK**, and then click **OK** again. See Figure 2-16.

TAKE NOTE*

You cannot encrypt a file with EFS while compressing a file with NTFS. You can only do one or the other.

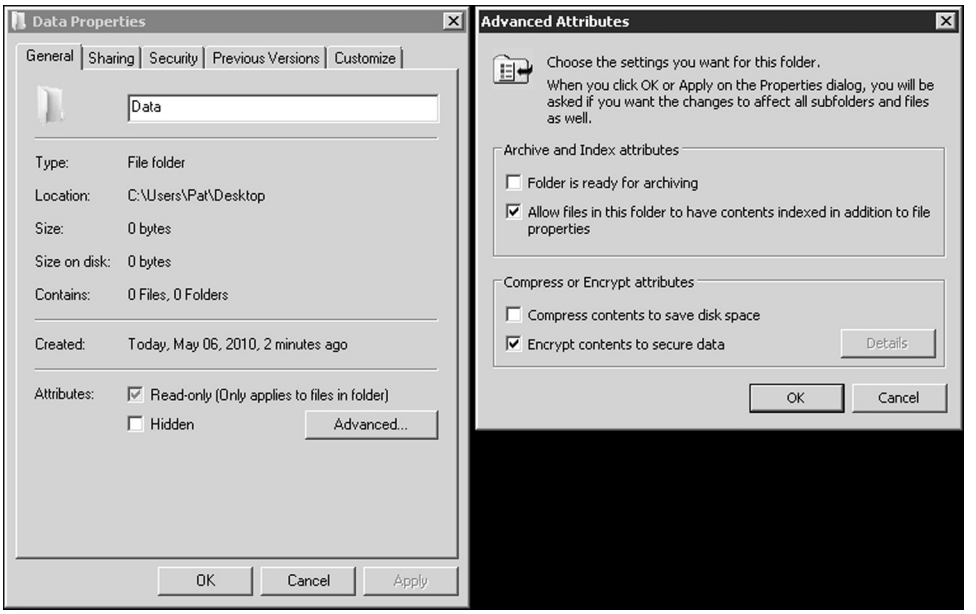

Figure 2-16

Encrypting data with EFS

DECRYPT A FOLDER OR FILE

GET READY. To decrypt a folder or file, perform the following steps:

1. Right-click the folder or file you want to decrypt, then click **Properties**.
2. Click the **General** tab, and then click **Advanced**.
3. Clear the **Encrypt contents to secure data** check box, click **OK**, and then click **OK** again.

The first time you encrypt a folder or file, an encryption certificate is automatically created. If your certificate and key are lost or damaged and you don't have a backup, you won't be able to use the files that you have encrypted. Therefore, you should back up your encryption certificate.

BACK UP EFS CERTIFICATE

GET READY. To back up your EFS certificate, perform the following steps:

1. Execute **certmgr.msc**. If you are prompted for an administrator password or confirmation, type the password or provide confirmation.
2. In the left pane, double-click **Personal**.
3. Click **Certificates**.
4. In the main pane, click the certificate that lists **Encrypting File System** under **Intended Purposes**. If there is more than one EFS certificate, you should back up all of them.
5. Click the **Action** menu, Select **All Tasks**, and then click **Export**.
6. In the **Certificate Export** wizard, click **Next**, click **Yes, export the private key**, and then click **Next**.
7. Click **Personal Information Exchange**, and then click **Next**.
8. Type the password you want to use, confirm it, and then click **Next**. The export process will create a file to store the certificate.
9. Type a name for the file and the location (include the whole path) or instead click **Browse**, navigate to a location, type a filename, and then click **Save**.
10. Click **Next**, and then click **Finish**.

You should then place the certificate in a safe place.

If for some reason, a person leaves your organization and you cannot read his or her encrypted files, you can also set up recovery agents who can recover encrypted files for a domain.

ADD USERS AS RECOVERY AGENTS

GET READY. To add new users as recovery agents, these users must first have recovery certificates issued by the enterprise CA structure.

1. Open the **Active Directory Users and Computers** console.
2. Right-click the domain and select **Properties**.
3. Select the **Group Policy** tab.
4. Select the **Default Domain Policy** and click **Edit**.
5. Expand **Computer Configuration\Windows Settings\Security Settings\Public Key Policies\Encrypted Data Recovery Agents**.
6. Right-click **Encrypted Data Recovery Agents** and select **Add**.
7. Click **Next** to the **Add Recovery Agent Wizard**.
8. Click **Browse Directory**. Locate the user and click **OK**.
9. Click **Next**.
10. Click **Finish**.
11. Close the **Group Policy Editor**.

If you copy a file or folder, the new file or folder will automatically acquire the encryption attribute of the original drive or folder. If the file or folder is moved within the same volume, it will retain the original assigned encryption attribute. Thus, if it is encrypted, it will remain encrypted at the new location. When the file or folder is moved from one volume to another, it is copied to the new location and then deleted from the old location. Therefore, the moved folder and files are new to the volume and acquire the new encryption attribute.

Encrypting Disks in Windows

Unlike EFS, BitLocker allows you to encrypt entire disks. Therefore, if a drive or laptop is stolen, the data is still encrypted, even if the thief installs it on another system of which he or she is an administrator.

TAKE NOTE*

BitLocker is a feature of Windows 7 Enterprise and Windows 7 Ultimate. It is not supported on other editions of Windows 7.

BitLocker Drive Encryption is the feature in the Windows 7 Ultimate and Enterprise editions that makes use of a computer's Trusted Platform Module (TPM). A TPM is a microchip built into a computer that is used to store cryptographic information, such as encryption keys. Information stored on the TPM can be more secure from external software attacks and physical theft. For instance, BitLocker Drive Encryption can use a TPM to validate the integrity of a computer's boot manager and boot files at startup, as well as to guarantee that a computer's hard disk has not been tampered with while the operating system was offline. BitLocker Drive Encryption also stores measurements of core operating system files in the TPM.

The system requirements of BitLocker are as follows:

- Because BitLocker stores its own encryption and decryption key in a hardware device that is separate from your hard disk, you must have one of the following:
  - A computer with Trusted Platform Module (TPM). If your computer was manufactured with TPM version 1.2 or higher, BitLocker will store its key in the TPM.
  - A removable USB memory device, such as a USB flash drive. If your computer doesn't have TPM version 1.2 or higher, BitLocker will store its key on the flash drive.
- Your computer must also have at least two partitions: a system partition (which contains the files needed to start your computer and must be at least 200 MB) and an operating system partition (which contains Windows). The operating system partition will be encrypted, and the system partition will remain unencrypted so your computer can start. If your computer doesn't have two partitions, BitLocker will create them for you. Both partitions must be formatted with the NTFS file system.
- In addition, your computer must have a BIOS that is compatible with TPM and supports USB devices during computer startup. If this is not the case, you will need to update the BIOS before using BitLocker.

BitLocker has five operational modes, which define the steps involved in the system boot process. These modes, in descending order from most to least secure, are as follows:

- **TPM + startup PIN + startup key:** The system stores the BitLocker volume encryption key on the TPM chip, but an administrator must supply a personal identification number (PIN) and insert a USB flash drive containing a startup key before the system can unlock the BitLocker volume and complete the system boot sequence.
- **TPM + startup key:** The system stores the BitLocker volume encryption key on the TPM chip, but an administrator must insert a USB flash drive containing a startup key before the system can unlock the BitLocker volume and complete the system boot sequence.
- **TPM + startup PIN:** The system stores the BitLocker volume encryption key on the TPM chip, but an administrator must supply a PIN before the system can unlock the BitLocker volume and complete the system boot sequence.
- **Startup key only:** The BitLocker configuration process stores a startup key on a USB flash drive, which the administrator must insert each time the system boots. This mode does not require the server to have a TPM chip, but it must have a system BIOS that supports access to the USB flash drive before the operating system loads.
- **TPM only:** The system stores the BitLocker volume encryption key on the TPM chip, and it accesses this key automatically when the chip has determined that the boot environment is unmodified. This unlocks the protected volume and the computer continues

to boot. Therefore, no administrative interaction is required during the system boot sequence.

When you enable BitLocker using the BitLocker Drive Encryption control panel, you can select the TPM + startup key, TPM + startup PIN, or TPM only options. To use the TPM + startup PIN + startup key option, you must first configure the *Require additional authentication at startup* Group Policy setting, found in the Computer Configuration\Policies\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives container.

ENABLING BITLOCKER

BitLocker is not enabled by default. If you don't know if your laptop comes with TPM, you should first verify that you have TPM. You will then turn on BitLocker for the volume that you wish to encrypt.

DETERMINE WHETHER YOU HAVE TPM

GET READY. To find out whether your computer has Trusted Platform Module (TPM) security hardware, perform the following steps:

1. Open the **Control Panel**, click **System and Security**, and click **BitLocker Drive Encryption**.
2. In the left pane, click **TPM Administration**. If you are prompted for an administrator password or confirmation, type the password or provide confirmation.

The TPM Management on Local Computer snap-in tells you whether your computer has the TPM security hardware. See Figure 2-17. If your computer doesn't have it, you'll need a removable USB memory device to turn on BitLocker and store the BitLocker startup key that you'll need whenever you start your computer.

Figure 2-17

TMP management console

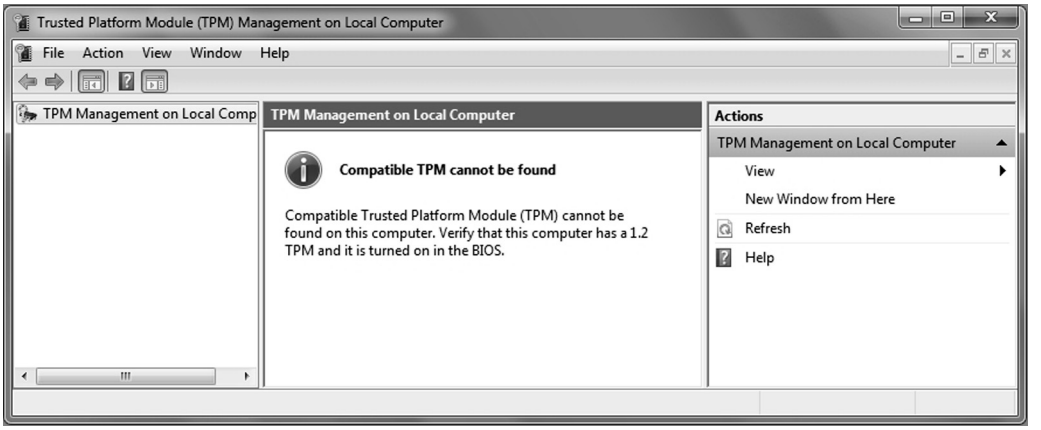

TURN ON BITLOCKER

GET READY. Log on to Windows 7 using an account with administrative privileges. Then, perform the following steps:

1. Click **Start**, then click **Control Panel > System and Security > BitLocker Drive Encryption**. The BitLocker Drive Encryption control panel appears.
2. Click **Turn on BitLocker** for your hard disk drives. The *Set BitLocker startup preferences* page appears. See Figure 2-18.

Figure 2-18
Turning on BitLocker

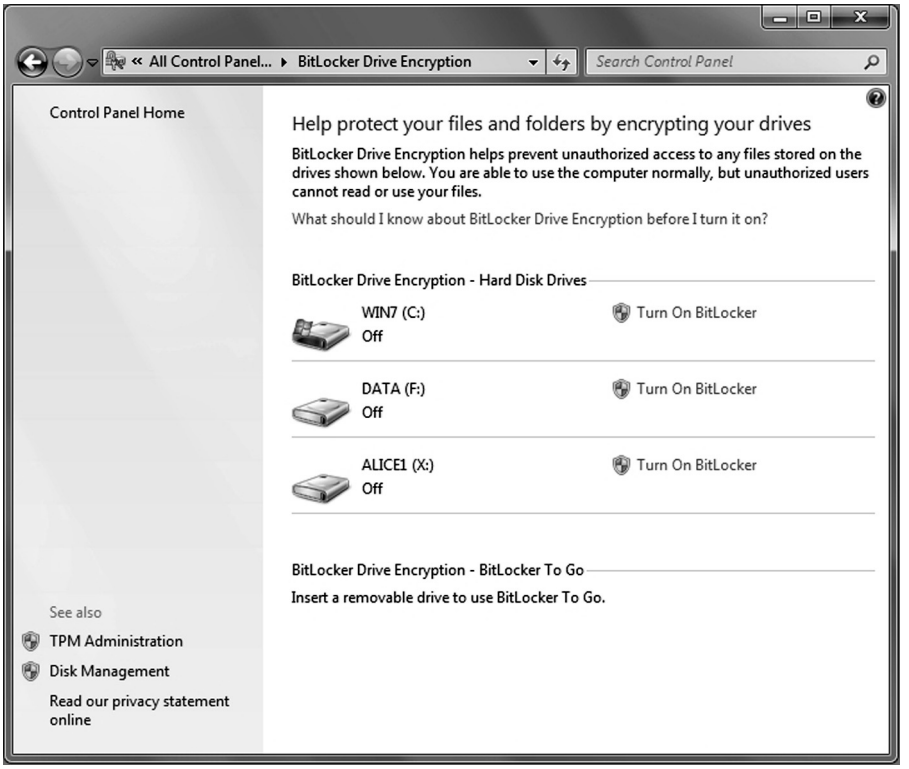

MORE INFORMATION

If your computer has a TPM chip, Windows 7 provides a Trusted Platform Module (TPM) management console that you can use to change the chip's password and modify its properties.

3. Click **Require a Startup key at every startup**. A *Save your Startup key* page appears.
4. Insert a USB flash drive into a USB port and click **Save**. The *How do you want to store your recovery key?* page appears.
5. Select one of the options to save your recovery key and click **Next**. The *Are you ready to encrypt this drive?* page appears.
6. Click **Continue**. The wizard performs a system check and then restarts the computer.
7. Log on to the computer. Windows 7 proceeds to encrypt the disk.

Once the encryption process is complete, you can open the BitLocker Drive Encryption control panel to ensure that the volume is encrypted or to turn off BitLocker when performing a BIOS upgrade or other system maintenance.

The BitLocker control panel applet enables you to recover the encryption key and recovery password at will. You should carefully consider how to store this information, because it will allow access to the encrypted data. It is also possible to escrow this information into Active Directory.

USING DATA RECOVERY AGENTS AND BITLOCKER

If for some reason, a user loses the startup key and/or startup PIN needed to boot a system with BitLocker, that user can supply the recovery key created during the BitLocker configuration process and gain access to the system. However, if the user loses the recovery key, you can use a data recovery agent designated with Active Directory to recover the data on the drive.

A data recovery agent (DRA) is a user account that an administrator has authorized to recover BitLocker drives for an entire organization with a digital certificate on a smart card. In most cases, administrators of Active Directory Domain Services (AD DS) networks use DRAs to ensure access to their BitLocker-protected systems and to avoid having to maintain large numbers of individual keys and PINs.

To create a DRA, you must first add the user account you want to designate to the Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies\BitLocker Drive Encryption container in a GPO or to the system's Local Security Policy. Then, you must configure the Provide The Unique Identifiers For Your Organization policy setting in the Computer Configuration\Policies\Administrative Templates\Windows Components\ BitLocker Drive Encryption container with unique identification fields for your BitLocker drives.

Finally, you must enable DRA recovery for each type of BitLocker resource you want to recover by configuring the following policies:

- Choose How BitLocker-Protected Operating System Drives Can Be Recovered
- Choose How BitLocker-Protected Fixed Drives Can Be Recovered
- Choose How BitLocker-Protected Removable Drives Can Be Recovered

These policies enable you to specify how BitLocker systems should store their recovery information, and they also enable you to store that information in the AD DS database.

USING BITLOCKER TO GO

BitLocker To Go is a new feature in Windows 7 that enables users to encrypt removable USB devices, such as flash drives and external hard disks. Although BitLocker has always supported the encryption of removable drives, BitLocker To Go allows you to use the encrypted device on other computers without having to perform an involved recovery process. Because the system is not using the removable drive as a boot device, a TPM chip is not required.

To use BitLocker To Go, simply insert the removable drive and open the BitLocker Drive Encryption control panel. The device appears in the interface, with a *Turn on BitLocker* link just like that of the computer's hard disk drive.

■ Introducting IPsec

THE BOTTOM LINE

IP Security, more commonly known as ***IPsec***, is a suite of protocols that provides a mechanism for data integrity, authentication, and privacy for the Internet Protocol. It is used to protect data that is sent between hosts on a network by creating secure electronic tunnels between two machines or devices. IPsec can be used for remote access, VPN server connections, LAN connections, or WAN connections.

IPsec ensures that data cannot be viewed or modified by unauthorized users while it is being sent to its destination. Before data is sent between two hosts, the source computer encrypts the information by encapsulating each data packet in a new packet that contains the information necessary to set up, maintain, and tear down the tunnel when it is no longer needed. The data is then decrypted at the destination computer.

There are a couple of modes and a couple of protocols available in IPsec depending on whether they are implemented by the end hosts (such as the server) or implemented on the routers and the desired level of security. In particular, IPsec can be used in one of two modes:

- **Transport mode:** Used to secure end-to-end communications, such as between a client and a server.
- **Tunnel mode:** Used for server-to-server or server-to-gateway configurations. The tunnel is the path a packet takes from the source computer to the destination computer. This way, any IP packets sent between the two hosts or between the two subnets, depending on the configuration, are secured.

In addition, the two IPsec protocols are as follows:

- **Encapsulating Security Payload (ESP):** Provides confidentiality, authentication, integrity, and antireplay for the IP payload only, not the entire packet. ESP operates directly on top of IP.
- **Authentication Header (AH):** Provides authentication, integrity, and antireplay for the entire packet (both the IP header and the data payload carried in the packet). It does not provide confidentiality, which means that it does not encrypt the payload. The data is readable but protected from modification. Some fields that are allowed to change in transit are excluded because they need to be modified as they are relayed from router to router. AH operates directly on top of IP.

ESP and AH can be combined to provide authentication, integrity, and antireplay for the entire packet (both the IP header and the data payload carried in the packet), as well as confidentiality for the payload.

Although AH and ESP provide the means to protect data from tampering, preventing eavesdropping and verifying the origin of the data, it is the Internet Key Exchange (IKE) that defines the method for the secure exchange of the initial encryption keys between the two endpoints. IKE allows nodes to agree on authentication methods, encryption methods, what keys to use, and the lifespan of the keys.

The information negotiated by IKE is stored in a Security Association (SA). An SA is like a contract laying out the rules of the VPN connection for the duration of the SA. Each SA is assigned a 32-bit number that, when used in conjunction with the destination IP address, uniquely identifies the SA. This number is called the Security Parameters Index (SPI).

IPsec can be used with Windows in various ways. To enable IPsec communications for a Windows Server 2008 computer, you would create group policies and assign them to individual computers or groups of computers. You could also use the Windows Firewall with advanced security.

Encrypting with VPN Technology

Today, it is common for organizations to use remote access server (RAS), which enables users to connect remotely via various protocols and connection types. By connecting to RAS over the Internet, users can connect to their organization's network so that they can access data files, read email, and access other applications just as if they were sitting at work. However, because the Internet is considered an insecure medium, you must use data encryption when setting up these types of connections.

A ***virtual private network (VPN)*** links two computers through a wide-area network such as the Internet. To keep the connection secure, the data sent between the two computers is encapsulated and encrypted. In one scenario, a client connects to the RAS server to access internal resources from offsite. Another scenario is to connect one RAS server on one site or organization to another RAS server on another site or organization so that the site or organizations can communicate with each other.

The four types of tunneling protocols used with a VPN server/RAS server running on Windows Server 2008 R2 are as follows:

- **Point-to-Point Tunneling Protocol (PPTP):** A VPN protocol based on the legacy Point-to-Point protocol used with modems. Unfortunately, PPTP is easy to set up but uses weak encryption technology.
- **Layer 2 Tunneling Protocol (L2TP):** Used with IPsec to provide security. This the industry standard when setting up secure tunnels.

- **Secure Sockets Tunneling Protocol (SSTP):** Introduced with Windows Server 2008, which users the HTTPS protocol over TCP port 443 to pass traffic through firewalls and web proxies that might block PPTP and L2TP/IPsec.
- **Internet Key Exchange version 2 (IKEv2):** Uses IPsec for encryption while supporting VPN Reconnect (also called Mobility), which enables VPN connections to be maintained when a VPN client moves between wireless cells or switches and to automatically reestablish broken VPN connectivity. Different from L2TP with IPsec, IKEv2 client computers do not need to provide authentication through a machine certificate or a preshared key.

When using VPNs, Windows 7 and Windows Server 2008 support the following forms of authentication:

- **Password Authentication Protocol (PAP):** Uses plain text (unencrypted passwords). PAP is the least secure form of authentication and is not recommended.
- **Challenge Handshake Authentication Protocol (CHAP):** A challenge-response authentication method that uses the industry standard md5 hashing scheme to encrypt the response. CHAP was an industry standard for years and is still quite popular.
- **Microsoft CHAP version 2 (MS-CHAPv2):** Provides two-way authentication (mutual authentication). MS-CHAPv2 provides stronger security than CHAP.
- **Extensible Authentication Protocol Microsoft CHAP version 2 (EAP-MS-CHAPv2):** EAP is a universal authentication framework that allows third-party vendors to develop custom authentication schemes including retinal scans, voice recognition, fingerprint identifications, smart cards, Kerberos, and digital certificates. It also provides a mutual authentication method that supports password-based user or computer authentication.

CREATE A VPN TUNNEL

GET READY. To create a VPN tunnel on a computer running Windows 7 so that you can connect to a Remote Access Server, perform the following steps:

1. From **Control Panel**, select **Network and Internet** to access the **Network and Sharing Center**.
2. From the **Network and Sharing Center**, choose **Set up a new connection wizard**.
3. In the **Set Up a Connection or Network** page, choose **Connect to a workplace**.
4. In the **Connect to a Workplace** page, answer the question: **Do you want to use a connection that you already have?** Choose whether you want to create a new connection or use an existing connection.
5. On the next page, choose **Use my Internet connection (VPN)**.
6. On the next screen, either choose your VPN connection or specify the Internet address for the VPN server and a destination name. You can also specify the following options: **Use a Smart card for authentication**, **Allow other people to use this connection**, and **Don't connect now**, **just set up so I can connect later**.

Often, you may need additional configurations of your VPN connection, such as those specifying the type of protocol, which authentication protocol to use, and the type of encryption.

After the VPN connection is created and configured, to connect using the VPN, simply open the Network and Sharing Center and click Manage Network Connections. Then right-click your VPN connection and click the Connect button. See Figure 2-19.

Figure 2-19
VPN connection

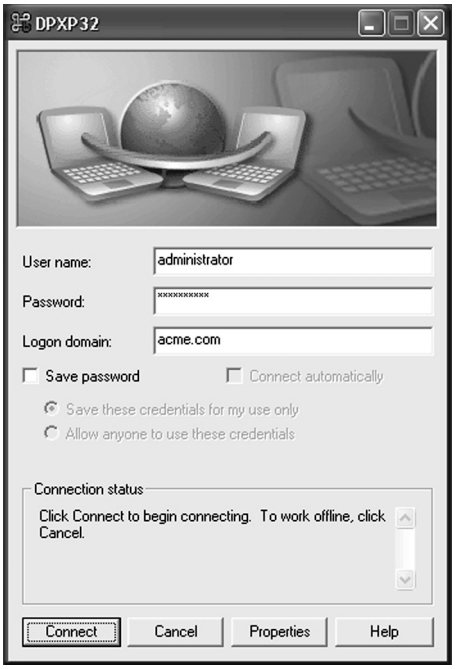

By default, when you connect to a VPN using the previous configuration, all web browsing and network traffic goes through the default gateway on the Remote Network unless you are communicating with local home computers. Having this option enabled helps protect the corporate network because all traffic will also go through firewalls and proxy servers, which helps prevent a network from being infected or compromised.

If you wish to route your browsing through your home Internet connection rather than through the corporate network, you can disable the "Use Default Gateway on Remote Network" option. When you disable this option, you are using what is known as split tunneling.

ENABLE SPLIT TUNNELING

GET READY. To enable split tunneling, perform the following steps:

1. Right-click a VPN connection and click **Properties**.
2. Click the **Networking** tab.
3. Double-click **Internet Protocol Version 4 (TCP/IPv4)**.
4. Click the **Advanced** button.
5. Deselect the **Use default gateway on remote network** option.

It can be a lot of work to configure multiple clients to connect to a remote access server. In fact, this task is often too complicated for computer novices, and it may be prone to errors. To help simplify administration of the VPN client into an easy-to-install executable, you could use the Connection Manager Administration Kit (CMAK). To install CMAK on Windows Server 2008, you must install it as a feature.

■ Using Auditing to Complete the Security Picture

THE BOTTOM LINE

As mentioned earlier, security can be divided into three areas. Authentication is used to prove the identity of a user, whereas authorization gives access to an authenticated user. To complete the security picture, however, you need to enable auditing so that you can have a record of which users have logged in and what resources those users accessed or tried to access.

CERTIFICATION READY
Can you explain why auditing is so important to security?
2.4

It is important that you protect your information and service resources from people who should not have access to them, while at the same time making those resources available to authorized users. Therefore, along with authentication and authorization, you should also enable auditing so that you can have a record of the following details:

- Who has successfully logged in
- Who has attempted to log in but failed
- Who has changed accounts in Active Directory
- Who has accessed or changed certain files
- Who has used a certain printer
- Who has restarted a system
- Who has made some system changes

Auditing is not enabled by default in Windows. To enable auditing, you must specify what types of system events to audit using group policies or the local security policy (Security Settings\Local Policies\Audit Policy). See Figure 2-20. Table 2-3 shows the basic audit events that are available in Windows Server 2003 and 2008. Windows Server 2008 also has additional options for more granular control. After you enable logging, you then open the Event Viewer security logs to view the logged security events. By default, these logs can only be seen and managed by the Administrators group.

Figure 2-20

Enabling auditing using group policies

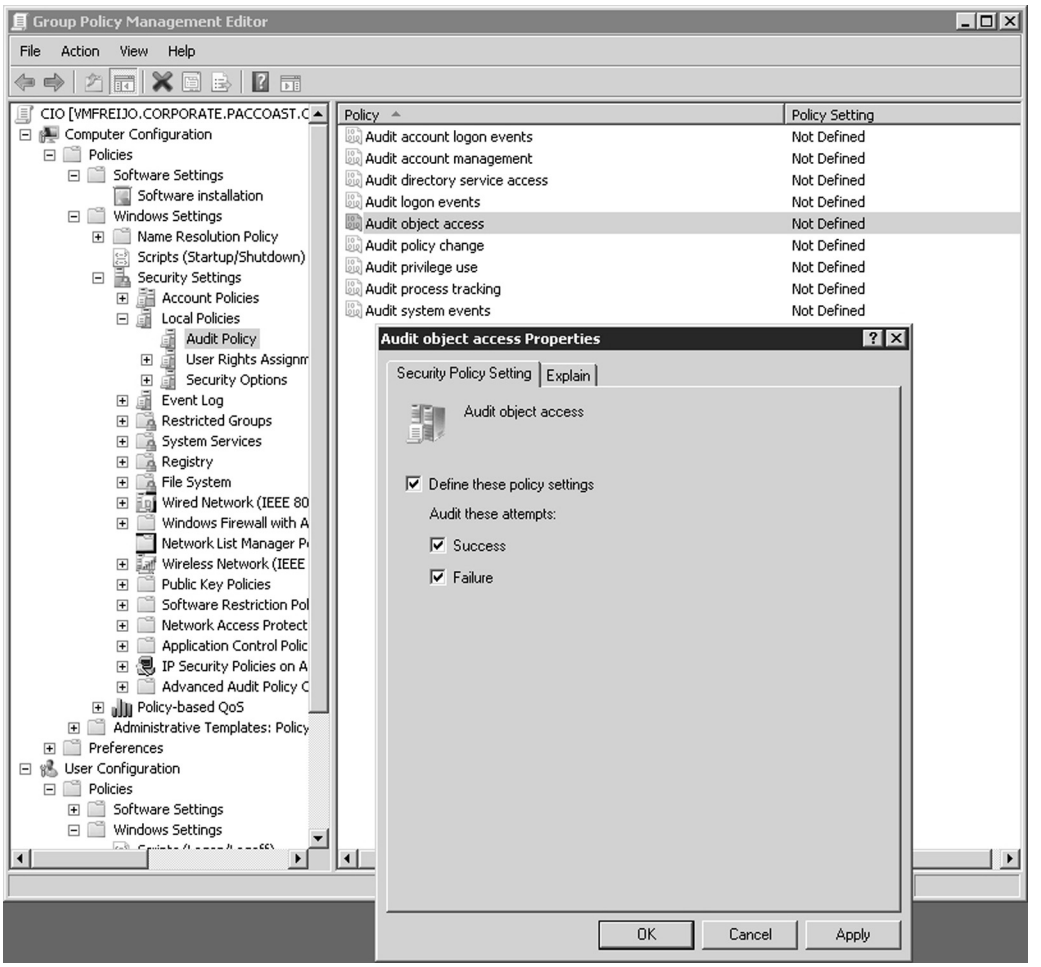

Table 2-3
Audit events

Event	Explanation
Account Logon	Determines whether the OS audits each time the computer validates an account's credentials, such as account login.
Account Management	Determines whether to audit each event of account management on a computer, including changing passwords and creating or deleting user accounts.
Directory Service Access	Determines whether the OS audits user attempts to access Active Directory objects.
Logon	Determines whether the OS audits each instance of a user attempting to log on or log off of his or her computer.
Object Access	Determines whether the OS audits user attempts to access non-Active Directory objects, including NTFS files, folders, and printers.
Policy Change	Determines whether the OS audits each instance in which users attempt to change user rights assignments, auditing policy, account policy, or trust policy.
Privilege Use	Determines whether to audit each instance in which a user exercises a user right.
Process Tracking	Determines whether the OS audits process-related events, such as process creation, process termination, handle duplication, and indirect object access. This is usually used for troubleshooting.
System	Determines whether the OS audits changes to the system time, system start up or shut down, attempts to load extensible authentication components, losses of auditing events due to auditing system failure, and security logs exceeding a configurable warning threshold level.

Auditing NTFS files, folders, and printers is a two-step process. You must first enable Object Access using group policies. Then you must specify which files, folders, or printers you want to audit. After enabling logging, you can open the Event Viewer security logs to view the security events.

Because Windows is only part of what makes up a network, you also need to look at other areas to audit. For example, for Microsoft's web server IIS, you can enable logging of who visits each site. For Microsoft's Internet Security and Acceleration (ISA) and Microsoft's Threat Management Gateway (TMG) servers, you can choose to log who accesses your network through a VPN or what is accessed through the firewall. Also, if you have Cisco routers and firewalls, you should enable auditing so that if someone reconfigures the router and firewall, you have a record of it.

If you need to audit non-Microsoft products, you may need to use Syslog. ***Syslog*** is a standard for logging program messages that can be accessed by devices that would not otherwise have a method for communications. Cisco firewalls and routers, computers running Linux and UNIX, and many printers can use Syslog. It can be employed for computer system management and security auditing, as well as for generalized information, analysis, and debugging messages.

After you decide what you are going to audit, you need to decide where you are going to keep the logs. You need to choose a server or device that has enough storage to hold the logs for the time required by your organization. You should also limit access to this storage area only to essential people. You should also consider backing up these logs and keeping the backups as long as required by your organization.

Lastly, if your organization is large enough or you have high security standards, you should consider having different people as administrators and different people as auditors. By having isolation of duties, the auditors can make sure that the administrators are doing what they are supposed to be doing and more importantly to make sure they are not doing what they are not supposed to be doing.

Finally, you should make sure that you have a change management system and a ticket system. A change management system will record what changes are made. It gives the IT department a method to review changes before they are implemented so that if these changes cause problems with a system, they can be evaluated. In addition, if a problem does occur, this system provides a single list of all of the changes made to your environment.

In comparison, a ticket system gives you a record of all problems and requests by users. By having a ticket system, you can determine what your most common problems are and identify trends.

AUDIT FILES AND FOLDERS

GET READY. Assuming that object auditing has been enabled, to audit files and folders, perform the following steps:

1. Open **Windows Explorer**.
2. Right-click the file or folder that you want to audit, click **Properties**, and then click the **Security** tab.
3. Click **Edit**, and then click **Advanced**.
4. In the **Advanced Security Settings for <object>** dialog box, click the **Auditing** tab.
5. Do one of the following:
   - To set up auditing for a new user or group, click **Add**. In **Enter the object name to select**, type the name of the user or group that you want, and then click **OK**.
   - To remove auditing for an existing group or user, click the group name or username, click **Remove**, click **OK**, and then skip the rest of this procedure.
   - To view or change auditing for an existing group or user, click the name of the group or user, then click **Edit**.
6. In the **Apply onto** box, click the location where you want auditing to take place.
7. In the **Access** box, indicate what actions you want to audit by selecting the appropriate check boxes:
   - To audit successful events, select the **Successful** check box.
   - To stop auditing successful events, clear the **Successful** check box.
   - To audit unsuccessful events, select the **Failed** check box.
   - To stop auditing unsuccessful events, clear the **Failed** check box.
   - To stop auditing all events, click **Clear All**.
8. If you want to prevent subsequent files and subfolders of the original object from inheriting these audit entries, select the **Apply these auditing entries to objects and/or containers within this container only** check box.
9. Click **OK** to close the Advanced Security Settings dialog box.
10. Click **OK** to close the Properties dialog box.

SKILL SUMMARY

In this lesson you learned:

- AAA (authentication, authorization, and accounting) is a model for access control.
- Authentication is the process of identifying an individual.
- After a user is authenticated, he or she can access network resources based on his or her authorization. Authorization is the process of giving individuals access to system objects based on their identity.
- Accounting, also known as auditing, is the process of keeping track of a user's activity when accessing network resources, including the amount of time spent in the network, the services accessed while in the network, and the amount of data transferred during the session.
- Nonrepudiation prevents one party from denying the actions it has carried out.
- Users can authenticate using what they know, what they own or possess, and/or what they are.
- When you use two or more methods to authenticate a user, you are implementing a multifactor authentication system.
- The most common method of authentication with computers and networks is the password.
- A password is a secret series of characters that enables a user to access a file, computer, or program.
- To hack a password, users will try obvious passwords, brute force attacks, and dictionary attacks.
- For increased security, you need to choose a password that nobody can guess. Therefore, your password should be long enough, and it should be considered strong or complex.
- A personal identification number (PIN) is a secret numeric password shared between a user and a system that can be used to authenticate the user to the system.
- A digital certificate is an electronic document that contains an identity, such as a user or organization, and a corresponding public key.
- A smart card is a pocket-sized card with embedded integrated circuits that consist of nonvolatile memory storage components and perhaps dedicated security logic.
- A smart card can contain digital certificates to prove the identity of the person carrying the card, and it may also contain permissions and access information.
- Biometrics is an authentication method that identifies and recognizes people based on physical traits, such as fingerprints, face recognition, iris recognition, retinal scans, and voice recognition.
- Because administrators have full access to computers and networks, you should use a standard nonadministrator account to perform most tasks.
- Active Directory is a technology created by Microsoft that provides a variety of network services, including LDAP, Kerberos-based and single sign-on authentication, DNS-based naming and other network information, and a central location for network administration and delegation of authority.
- Kerberos is the default computer network authentication protocol. It allows hosts to prove their identity over a nonsecure network in a secure manner.
- Single sign-on (SSO) allows you to log on once and access multiple related, but independent, software systems without having to log in again.
- A user account enables a user to log on to a computer and domain.
- Local user accounts are stored in the Security Account Manager (SAM) database on the local computer.
- Groups are used to group users and computers together so that when you assign rights and permissions, you can assign them to the entire group rather than to each user individually.
- A right authorizes a user to perform certain actions on a computer, such as logging on to a system interactively or backing up files and directories on a system.

- A permission defines the type of access granted to an object or object attribute.
- Explicit permissions are permissions granted directly to a file or folder.
- Inherited permissions are permissions that are granted to a folder (parent object or container) and then flow into the folder's child objects (subfolders or files inside the parent folder).
- The owner of an object controls how permissions are set on the object and to whom permissions are granted.
- Encryption is the process of converting data into a format that cannot be read by another user. Once a user has encrypted a file, that file automatically remains encrypted when it is stored on disk.
- Decryption is the process of converting data from an encrypted format back to its original format.
- Encryption algorithms can be divided into three classes: symmetric, asymmetric, and hash function.
- Symmetric encryption uses a single key to encrypt and decrypt data. Therefore, it is also referred to as secret-key, single-key, shared-key, and private-key encryption.
- Asymmetric key encryption, also known as public key cryptography, uses two mathematically related keys. One key is used to encrypt the data, while the second is used to decrypt it.
- Different from symmetric and asymmetric algorithms, a hash function is meant as a one-way encryption. That means that after information has been encrypted, it cannot be decrypted.
- Public key infrastructure (PKI) is a system consisting of hardware, software, policies, and procedures that create, manage, distribute, use, store, and revoke digital certificates.
- The most common digital certificate is the X.509 version 3.
- The certificate chain, also known as the certification path, is a list of certificates used to authenticate an entity. It begins with the certificate of the entity and ends with the root CA certificate.
- A digital signature is a mathematical scheme that is used to demonstrate the authenticity of a digital message or document It is also used to prove that the message or document has not been modified.
- When you need to transmit private data over the Internet, you should use SSL over HTTPS (https) to encrypt the data that you're sending. URLs that require an SSL connection start with https: instead of http:.
- IP Security, more commonly known as IPsec, is a suite of protocols that provides a mechanism for data integrity, authentication, and privacy for the Internet Protocol.
- A virtual private network (VPN) links two computers through a wide-area network, such as the Internet.
- Syslog is a standard for logging program messages that can be accessed by devices that would not otherwise have a method for communication.

■ Knowledge Assessment

Multiple Choice

Circle the letter that corresponds to the best answer.

1. Which of the following is not a method for authentication?
   a. Something the user knows
   b. Something the user owns or possesses
   c. Encryption
   d. Something the user is

2. Which of the following is not a biometric device?
   a. Password reader
   b. Retinal scanner
   c. Fingerprint scanner
   d. Face scanner

3. Which of the following services is used for centralized authentication, authorization, and accounting?
   a. VPN
   b. PGP
   c. RADIUS
   d. PKI

4. What is the primary authentication method used on Microsoft Active Directory?
   a. LDAP
   b. Kerberos
   c. NTLAN
   d. SSO

5. The master time keeper and master for password changes in an Active Directory domain is
   a. PDC Emulator
   b. RID
   c. Infrastructure master
   d. Schema master

6. Local user accounts are found in
   a. Active Directory
   b. Registry
   c. SAM
   d. LDAP

7. A(n) _____ authorizes a user to perform certain actions on a computer.
   a. Permission
   b. Encryption algorithm
   c. Authentication protocol
   d. Right

8. Which of the following file systems offers the best security?
   a. FAT
   b. FAT32
   c. NTFS
   d. EFS

9. Which NTFS permission is needed to change attributes and permissions?
   a. Full Control
   b. Modify
   c. Read and Execute
   d. Write

10. Which type of permission is granted directly to a file or folder?
    a. Explicit
    b. Inherited
    c. Effective
    d. Share

11. If you copy a file or folder to a new volume, what permissions will that file or folder have?
    a. The same permissions that it had before.
    b. The same permissions as the target folder.
    c. The same permissions as the source folder.
    d. No permissions at all.
12. Which of the following uses an ACL?
    a. NTFS folder
    b. Active Directory user
    c. Registry key
    d. Login rights
13. Which type of key has one key for encryption and a different key for decryption?
    a. Symmetric
    b. Asymmetric
    c. Hash function
    d. PKI
14. Which infrastructure is used to assign and validate digital certificates?
    a. Asymmetric algorithm
    b. Active Directory
    c. PKI
    d. VPN
15. Which technology is used to encrypt an individual file on an NTFS volume?
    a. BitLocker
    b. BitLocker To Go
    c. PPTP
    d. EFS

Fill in the Blank

Complete the following sentences by writing the correct word or words in the blanks provided.

1. A(n) _____ is a secret numeric password shared between a user and a system that can be used to authenticate the user to the system.
2. A pocket-sized card with embedded integrated circuits that is used for authentication is known as a(n) _____.
3. A device that may give you a second password to log in to a system is a(n) _____.
4. The _____ holds a copy of the centralized database used in Active Directory.
5. By default, your computer clock should not be off more than _____ minutes or you might have problems with Kerberos authentication.
6. A(n) _____ defines the type of access over an object or the properties of an object such as an NTFS file or printer.
7. _____ permissions flow from a parent object to a child object.
8. When you cannot access a folder because someone removed the permissions so that no one can access it, you must take _____ of the folder.
9. The centralized database that holds most of the Windows configuration is known as the _____.
10. To track a user's activities in Windows, you need to enable _____.

Competency Assessment

Scenario 2-1: Understanding the Disadvantages of Biometrics

You are the IT administrator for the Contoso Corporation. Your CIO wants you to investigate the possible use of biometrics for security purposes. The CIO understands what biometrics is and how this technology can be used, but he does not understand the potential disadvantages of using biometrics. What should you tell him?

Scenario 2-2: Limiting Auditing

You are the IT administrator for the Contoso Corporation. Your CIO needs to know when a particular user accessed a certain folder. However, this information is not available because auditing was not enabled. To ensure that this does not happen again in the future, the CIO asks you to enable auditing for everything. How should you respond?

Proficiency Assessment

Scenario 2-3: Looking at NTFS Permissions

Log in as an administrator on a computer running Windows 7 or Windows Server 2008. Create a group called Managers on your computer. Now, create a user account called JSmith and assign it to the Managers group. Next, create another user account called JHamid. Create a folder called SharedTest, and create a text file called test.txt in the SharedTest Folder. Share the folder. Assign Allow Full Control to Everyone. Assign Read and Execute to the Managers group. Log in as JHamid and try to access the \\localhost\SharedTest folder. Then, log in as JSmith and try access the \\localhost\SharedTest folder.

Scenario 2-4: Looking at EFS

Add JHamid to the Managers group you established in the previous exercise. Now, log in as JSmith and encrypt the test.txt file with EFS. Finally, log in as JHamid and try to access the test.txt file.

Workplace Ready

Planning and Maintaining Security

When considering security, you need to look at the entire picture. Security must be planned for from the beginning. Therefore, you need to define what your security goals are, what impact they will have on current access and network applications, and how security measures will affect users. Then, after such measures have been implemented, you must maintain them by constantly monitoring the security of the system, making changes as needed, patching security holes, and constantly reviewing the security logs.