Windows 2000 Operating Systems

This overview begins by taking a look at the Microsoft Windows 2000 operating system family. The operating systems that make up this family are:

- Windows 2000 Professional
- Windows 2000 Server
- Windows 2000 Advanced Server
- Windows 2000 Datacenter Server

These four operating systems share a common user interface, share many common features and utilities, and are all 32-bit operating systems. In fact, all of these operating systems use the same kernel, which is based on Windows NT technology.

Although based on the same kernel, each of the four operating systems that make up the Windows 2000 operating system family is optimized for use in a specific environment.

The following section explores some of the new common features shared by the four Windows 2000 operating systems.

Windows 2000 Professional

Microsoft Windows 2000 Professional is a 32-bit operating system that is optimized for use on desktop computers. Windows 2000 Professional picks up where Windows NT Workstation left off. It contains not only the features and functionality of Windows NT Workstation, but also the best features of Windows 98.

Windows 2000 Professional is typically not a good choice of operating system for a server in a business environment, because it supports only ten concurrent connections from other computers.

Hardware Requirements

As with all new versions of operating systems, Windows 2000 Professional requires significantly more hardware resources than did either of its predecessors — Windows NT Workstation or Windows 98. The minimum hardware required to successfully install and run Windows 2000 Professional on an Intel-based computer includes:

- A Pentium/133MHz processor
- 32MB of RAM (64MB are recommended)
- 650MB of free hard disk space

In order to ensure operational success, all hardware should be on the Windows 2000 *Hardware Compatibility List* (HCL) that is shipped with the product and is also posted on Microsoft's Web site.

Application Support

Windows 2000 Professional supports most MS-DOS–based applications, most 16-bit and 32-bit Windows-based applications, POSIX 1.*x* applications, and most OS/2 1.*x* applications. Specifically, Windows 2000 Professional supports many Windows 95/Windows 98 applications that were not supported by Windows NT Workstation 4.0. Windows 2000 Professional does not support applications that require direct hardware access (bypassing the Hardware Abstraction Layer [HAL]) because this

1. Explain Multithreading, Multiprocessing, Multitasking in Windows 2000 Professional.

could compromise Windows 2000 Professional's security. It also does not support software applications that require an MS-DOS terminate-and-stay-resident (TSR) program or a virtual device driver.

I'll discuss the various application environments supported by Windows 2000 in more detail a bit later in this chapter.

Multiprocessing, Multithreading, and Multitasking

Windows 2000 Professional supports symmetric multiprocessing with up to two processors. *Multiprocessing* refers to the capability of an operating system to use more than one processor in a single computer simultaneously. *Symmetric multiprocessing* is a type of multiprocessing in which system processes and applications can be run on any available processor. This is the most efficient form of multiprocessing currently available, because it does not tie a particular process or application to a specific, assigned processor.

Windows 2000 Professional also supports multithreading and preemptive multitasking. A *thread* is the smallest unit of processing that can be scheduled by the Windows 2000 kernel. All applications require at least one thread. When an application has more than one thread, each thread can be executed independently of the others. This is referred to as *multithreading*. Individual threads within a single application can even be run on different processors in the same computer. In *preemptive multitasking*, the operating system allocates processor time between applications. Because Windows 2000 — not the application — allocates processor time between multiple applications, one application can be preempted by the operating system, and another application allowed to run. When multiple applications are alternately paused and then allocated processor time, they appear to run simultaneously to the user.

Security

Windows 2000 Professional supports a high level of security. User logon and authentication are required in order to use the operating system and in order to access local or network resources. Windows 2000 Professional supports a local user account database, and can also support either a Windows NT Server 4.0 domain user account database or user accounts from the Windows 2000 Active Directory.

Two other security features of Windows 2000 Professional are smart card support and Internet Protocol Security. A *smart card* is a security device that contains a unique, encrypted set of authentication credentials. When used in

2. Describe File Management in Windows 2000 Server.

conjunction with a smart card reader that has been installed on a computer, smart cards eliminate the need for users to transmit user names and passwords across the network when logging on. *Internet Protocol Security* (IPSec) encrypts TCP/IP traffic between two computers, thus preventing unauthorized users who capture network traffic from viewing or modifying sensitive data.

Windows 2000 Server

Microsoft Windows 2000 Server is a powerful 32-bit operating system that is optimized for network file, print, application, and Web servers. Windows 2000 Server is the next generation of Windows NT Server. It contains all of the features and functionality of Windows 2000 Professional, plus several additional features that make it the operating system of choice for most business server applications.

Hardware Requirements

The minimum hardware required to successfully install and run Windows 2000 Server on an Intel-based computer includes:

- A Pentium/133MHz processor
- 64MB of RAM (128MB are recommended)
- 950MB of free hard disk space (more disk space is required if the computer contains more than 64MB of RAM)

All hardware should be on the Windows 2000 HCL.

File Management

Windows 2000 Server supports two new file management tools, the Distributed file system (Dfs) and disk quotas.

The *Distributed file system* (Dfs) is a file system that enables an administrator to make shares that are stored on various servers on the network appear to users as though they are stored within a single share on a single server. The use of Dfs makes finding network resources easier for users, because users don't have to know which server physically contains the shared resource they are trying to access.

Disk quotas is a volume management tool that is enabled on a volume-by-volume basis. Once enabled, disk quotas automatically track disk space usage on a user-by-user basis, and prevent individual users from exceeding the disk space limitations that they have been assigned by administrators.

Disk quotas can also be used on Windows 2000 Professional computers, but it seems unlikely to me that they will be widely used on desktop client computers.

Application Support

Windows 2000 Server supports the same software applications as Windows 2000 Professional. In addition, Windows 2000 Server is optimized to support the Microsoft BackOffice suite of products, including SQL Server, Systems Management Server, Internet Information Server, Exchange Server, and SNA Server, as well as many third-party server-based applications.

Windows 2000 Server also supports *Terminal Services*. This application service, when run on a network server, enables users of client computers to remotely perform processor-intensive or network-intensive tasks from their client computers. The application runs on the server running Terminal Services, so the user can take advantage of the processing power and network connectivity of the server, while fully controlling the application from the client computer's keyboard and monitor.

Multiprocessing, Multithreading, and Multitasking

Like Windows 2000 Professional, Windows 2000 Server supports symmetric multiprocessing, but Windows 2000 Server accommodates up to four processors instead of only two. Also like Windows 2000 Professional, Windows 2000 Server supports multithreading and preemptive multitasking.

Security

Windows 2000 Server includes all of the security features of Windows 2000 Professional, and has additional security features of its own.

Windows 2000 Server supports a local user account database, and can also support either a Windows NT Server 4.0 domain user account database, or user accounts from the Windows 2000 Active Directory. In addition, Windows 2000 Server can be configured as a domain controller, which

contains a read/write copy of the Active Directory data store. *Active Directory* is a directory service that stores information about various types of network objects, including printers, shared folders, user accounts, and computers. These objects are placed in a hierarchical structure that can be organized to simplify administration. With Active Directory, users can gain access to any network resource (that the user has permissions to) with a single logon.

Windows 2000 Server also includes support for Remote Authentication Dial-In User Service (RADIUS). *RADIUS* is an industry standard authentication service that provides centralized management of user authentication and authorization for remote access servers.

Networking

Windows 2000 Server supports routing of the IP, IPX, and AppleTalk protocols over both LAN and WAN interfaces. Both the Routing Information Protocol (RIP) version 2 and the Open Shortest Path First (OSPF) routing protocols are supported for IP routing.

Another new networking feature of Windows 2000 Server is the support this operating system provides for asynchronous transfer mode (ATM) network adapter cards. ATM technology makes possible the simultaneous transport of voice, data, video, and images over the network.

Windows 2000 Advanced Server

Microsoft Windows 2000 Advanced Server is a powerful 32-bit operating system that is optimized for servers in an enterprise network environment. This operating system is often also a good intermediate choice for a heavily used server, such as a SQL server, when you need a more powerful hardware platform than Windows 2000 Server supports, but don't need the capabilities (or the added hardware and software expense) associated with Windows 2000 Datacenter Server.

Windows 2000 Advanced Server provides more scalability than Windows 2000 Server. Windows 2000 Advanced Server supports up to eight processors, and up to 8GB of RAM. Windows 2000 Server, on the other hand, only supports up to four processors and up to 4GB of RAM.

The minimum hardware requirements of Windows 2000 Advanced Server are virtually the same as those for Windows 2000 Server. As noted previously, however, Windows 2000 Advanced Server can support more processors and more RAM than Windows 2000 Server.

Windows 2000 Advanced Server includes all of the features of Windows 2000 Server. In addition, Windows 2000 Advanced Server includes Windows Clustering. A *cluster* is a group of computers that, from a client and application point of view, appear as a single computer. *Windows Clustering* is a technology which, when implemented on 2 to 32 Windows 2000 Advanced Server computers, provides two important features:

- **High availability:** This feature is important for mission-critical applications. In Windows Clustering, if a computer in the cluster that is running a critical application fails, another computer in the cluster will automatically start the application, and users will be seamlessly directed to the computer that takes over running the application.
- **Load balancing:** This feature refers to spreading utilization across multiple computers. For example, if a Web server experiences more utilization than a single computer can handle, it can be run on all of the computers in the cluster. Users will be seamlessly directed to the computer with the lowest utilization.

Windows Clustering is implemented on Windows 2000 Advanced Server by installing the Cluster Service.

Windows 2000 Datacenter Server

Microsoft Windows 2000 Datacenter Server is the most powerful of the Windows 2000 operating systems. Also a 32-bit operating system, Windows 2000 Datacenter Server is optimized for enterprise applications, such as extremely large databases and real-time online transaction processing, or other industrial applications that require phenomenal amounts of processor power.

Windows 2000 Datacenter Server provides more scalability than Windows 2000 Advanced Server. Windows 2000 Datacenter Server supports up to 32 processors, and up to 64GB of RAM. Windows 2000 Advanced

Server, on the other hand, only supports up to eight processors and up to 8GB of RAM.

The minimum hardware requirements of Windows 2000 Datacenter Server are the same as those for Windows 2000 Server. As noted previously, however, Windows 2000 Datacenter Server can support more processors and more RAM than either Windows 2000 Server or Windows 2000 Advanced Server.

The features of Windows 2000 Datacenter Server are identical to the features of Windows 2000 Advanced Server. The only advantage of Windows 2000 Datacenter Server is its capability to utilize more processors and more RAM.

Architecture of Windows 2000

An overview of Windows 2000 wouldn't be complete without discussing its architecture. If you develop a basic understanding of the operating system's architecture now, you'll have a framework on which to "hang" all of the concepts and facts presented throughout the rest of this book.

Windows 2000 uses a modular architecture. This means each component (or module) within the architecture has sole responsibility for the function it is designed to provide. In addition, no other module repeats the functions performed by another. Figure 1-3 illustrates the modular architecture of Microsoft Windows 2000. Notice that the operating system has two parts, or modes: user mode and kernel mode.

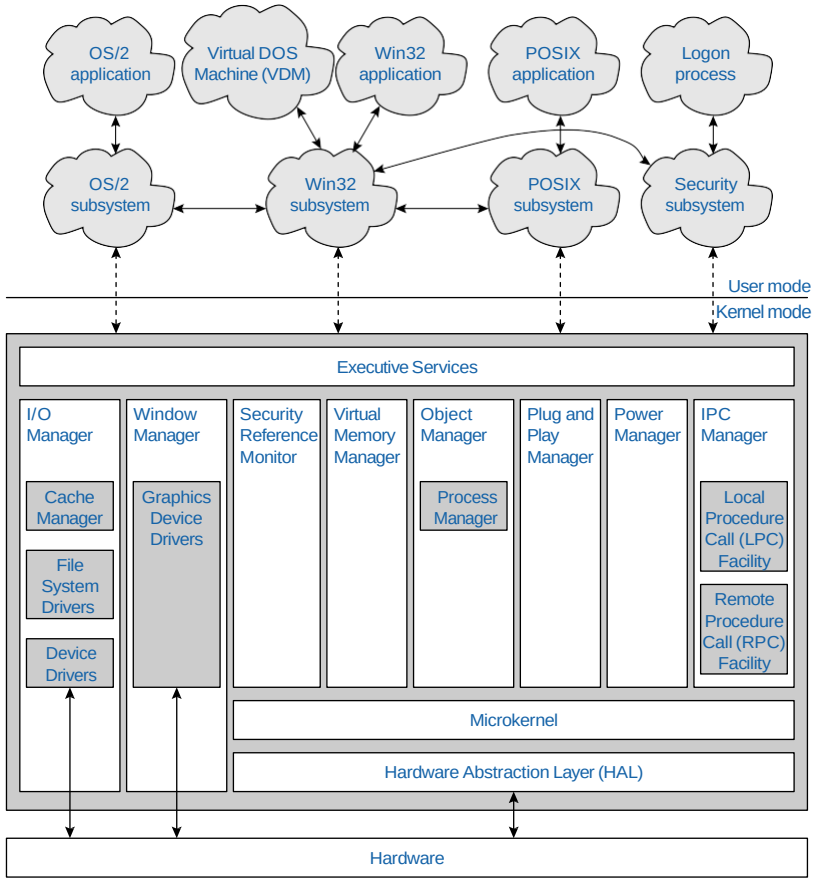

FIGURE 1-3 Microsoft Windows 2000 modular architecture

User Mode

Applications and their subsystems run in *user mode*. This mode is referred to as a less-privileged processor mode because it does not have direct access to hardware. User mode applications are limited to assigned memory address spaces and can't directly access other memory address spaces. User mode uses specific *application programming interfaces* (APIs) to request services from a kernel mode component.

The purpose of separating the applications in user mode from the hardware, of restricting the memory address spaces that applications can access, and of forcing the applications to run all requests for services through the kernel mode, is to protect against the possibility of an application crashing the system, and also to protect against unauthorized user access.

Examine Figure 1-3 again, and notice that there are four main subsystems in user mode: the OS/2 subsystem, the Win32 subsystem, the POSIX subsystem, and the Security subsystem.

The *OS/2 subsystem* is required to run OS/2 1.*x*–compatible applications. The OS/2 subsystem obtains its user interface and its screen functions from the Win32 subsystem, and requests Executive Services in kernel mode to perform all other functions for it. (Executive Services is covered in the next section of this chapter.)

The *Win32 subsystem* is the primary application subsystem. All 32-bit Windows applications run in this subsystem. The Win32 subsystem provides its own screen and keyboard functions, and requests Executive Services in kernel mode to perform all other functions for it. The Win32 subsystem also provides screen and keyboard functions for all of the other subsystems.

The *POSIX subsystem* is designed to run POSIX 1.*x*–compatible applications. It functions very much like the OS/2 subsystem. The POSIX subsystem uses the Win32 subsystem to provide all of its screen and graphical displays, and it requests Executive Services in kernel mode to perform all other functions for it.

Finally, the *Security subsystem*, which is also referred to as the Integral subsystem, supports the logon process. This subsystem also supports and provides the security for Active Directory. The Security subsystem obtains its user interface and its screen functions from the Win32 subsystem, and requests Executive Services in kernel mode to perform all other functions for it.

In addition to the four formal subsystems, a Virtual DOS Machine (VDM) is a feature of user mode. Its function is to run MS-DOS–based and Windows 3.*x*–based (all 16-bit) applications. Because the VDM is a Win32 application, all of its services, including screen and keyboard functions, are provided by the Win32 subsystem.

Kernel Mode

Kernel mode refers to a highly privileged mode of operation. It is called "highly privileged" because all code that runs in kernel mode can access the hardware directly, and can also directly access memory. A process running in kernel mode is not restricted to its own specific memory address space as is an application running in user mode.

The entire set of services that comprise kernel mode is called Executive Services (or sometimes the Windows NT Executive, or the Executive, for

short). Executive Services provide kernel mode services as requested by applications in user mode.

Notice how Figure 1-3 graphically presents the pieces of kernel mode. Kernel mode is made up of numerous components integral to the major Windows 2000 operating system functions.

The *Executive Services* component functions as an interface between user mode and kernel mode. Its purpose is to pass information between user mode subsystems and kernel mode components. In addition, Executive Services is responsible for the transfer of information and instructions between the various kernel mode components. Executive Services can be thought of as the "glue" that holds Windows 2000 together. As mentioned earlier, Executive Services is also called the Windows NT Executive, or the Executive, for short.

The *I/O Manager* is responsible for all input and output to disk storage subsystems. As it manages input and output, the I/O Manager also serves as a manager and supporter of communication between the various drivers. The I/O Manager can communicate directly with system hardware if it has the appropriate hardware device drivers. Subcomponents of the I/O Manager include a Cache Manager, File System Drivers, and Device Drivers.

Window Manager is responsible for providing the graphical user interface. Window Manager communicates directly with the graphics device drivers, which in turn communicate directly with the hardware. In the early days of Windows NT (versions 3.51 and earlier), Window Manager was an integral part of the Win32 subsystem in user mode. When Windows NT 4.0 came along, the developers moved Window Manager from user mode to kernel mode. This change enabled faster access to the graphics device drivers and eliminated the need for user mode applications to switch back and forth between kernel mode and user mode to make calls for graphics services. For these reasons, Window Manager continues to be a kernel mode component in Windows 2000.

There are six other kernel mode subsystems: the Security Reference Monitor, the Virtual Memory Manager, the Object Manager, the Plug and

3. Write a short note on Workgroups and Domains.

Play Manager, the Power Manager, and the IPC Manager. Each one of these subsystems communicates directly with the Microkernel.

The *Microkernel* is the very heart of the Windows 2000 operating system. It handles interrupts, schedules threads, and synchronizes processing activity. The Microkernel, in turn, communicates with the Hardware Abstraction Layer (HAL).

The *HAL* is designed to hide the varying characteristics of hardware so that all hardware platforms appear the same to the Microkernel. As a result, only the HAL, and not the entire Microkernel, needs to address each and every hardware platform. The HAL can communicate directly with the computer's hardware.

Now that you've been introduced to user mode and kernel mode, you're ready to move on to the last major architecture topic: the Windows 2000 memory model.

Workgroups, AND Domains

Before this overview of Microsoft Windows 2000 can be complete, it's important that you get good and comfortable with three key concepts: workgroups, domains, and Active Directory. First, I'll tackle workgroups and domains, and then I'll discuss Active Directory.

Workgroups and domains are two methods of grouping networked computers for common purposes. Computers and their users may be grouped based on common usage requirements or on departmental or geographical traits. For example, all the members of an accounting department or all the computers on the third floor of a building may be grouped together.

Workgroups

A *workgroup* is a logical grouping of networked computers in which one or more of the computers has one or more shared resources, such as a shared folder or a shared printer.

In a Windows 2000 (or Windows NT) workgroup environment, user account security is maintained individually at each separate computer through the use of a local user account database. Resources and administration are distributed throughout the computers that make up the workgroup. In a workgroup configuration there is no centrally maintained user accounts database, nor any centralized security. This means that a user must have a user account on each computer in the workgroup that

contains a shared resource that the user needs to access. Figure 1-4 illustrates how user account security is distributed throughout a workgroup environment. Notice that user account security is maintained individually at each sepa- rate computer in the workgroup.

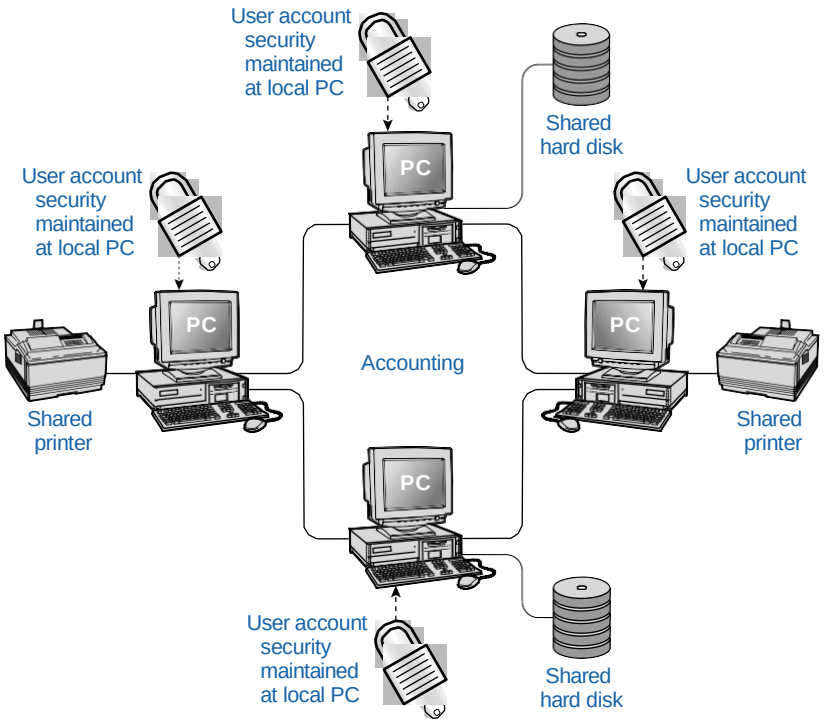

FIGURE 1-4 User account security in a workgroup environment

Typically, all of the computers in a workgroup run desktop operating systems, such as Windows 2000 Professional or Windows NT Workstation. Computers in a workgroup may also run Windows 95 or Windows 98, but these operating systems do not support a local user account database.

Workgroups are most often implemented in small networks where no centralized security or administration is desired. When a workgroup is used, the user of each computer controls access to the specific resources that are shared by that user's computer, and also maintains the computer's local user account database. It stands to reason, then, that the larger the workgroup, the more time and effort users must spend administering their local computers.

As a network becomes larger and more complex, administration and security become harder to manage. In these situations a domain (which is the subject of the next section) will most likely be used instead of a workgroup.

Domains

A *domain* is a logical grouping of networked computers in which one or more of the computers has one or more shared resources, such as a shared folder or a shared printer, *and* in which all of the computers share a common central domain directory database that contains user account security information.

One distinct advantage of using a domain, particularly on a large network, is that administration of user account security for the entire network can be managed from a centralized location. In a domain, a user has only one user account, which is stored in the domain directory database. This user account enables the user to access shared resources (that the user has permissions to access) located on any computer in the domain. Figure 1-5 illustrates how user account security is centralized in a domain environment. Note that all user account security is maintained by the domain controller.

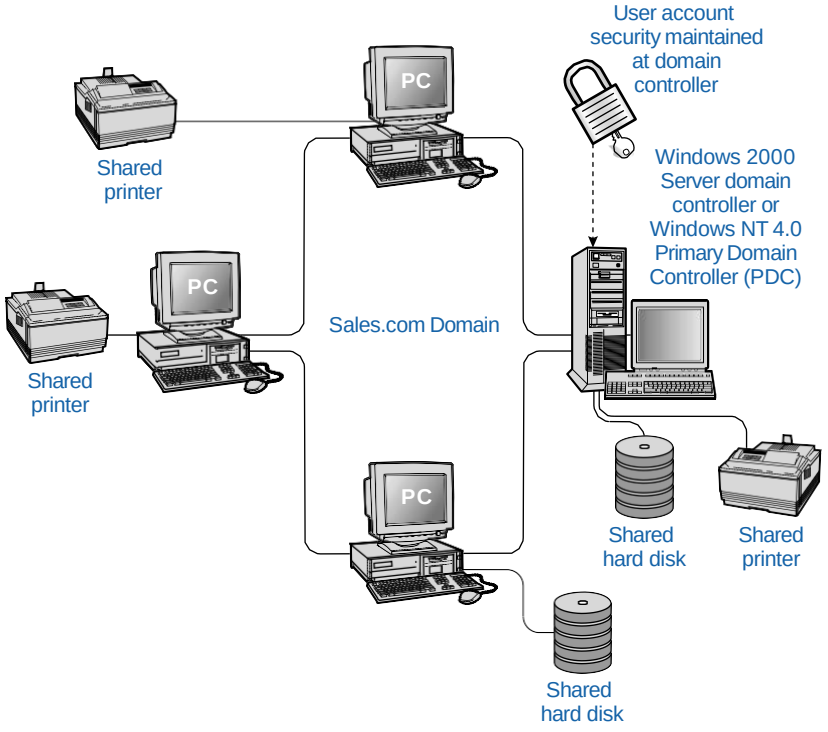

FIGURE 1-5 User account security in a domain environment

Domains are implemented differently in Windows 2000 than they are in Windows NT 4.0. The following sections explore the similarities and differences between Windows NT 4.0 domains and Windows 2000 domains.

Windows NT 4.0 Domains

In a Windows NT 4.0 domain, at least one of the networked computers is a server that runs Windows NT Server 4.0. This server is configured as a *primary domain controller* (PDC), which maintains the domain directory database. Typically, there is at least one additional server that also runs Windows NT Server. This additional server (or servers) is usually configured as a *backup domain controller* (BDC). The other computers on the network normally run a client operating system, such as Windows 95, Windows 98, or Windows NT Workstation. Resources, such as hard disks and printers, can be shared from any computer on the network.

Windows 2000 Domains

In a Windows 2000 domain, at least one of the networked computers is a server that runs Windows 2000 Server. This server is configured as a *domain controller*, which maintains the Active Directory data store. Typically, there is at least one additional server computer that also runs Windows 2000 Server. This additional computer is also usually configured as a domain controller, which contains a read/write copy of the Active Directory data store. The purpose of the additional server (or servers) is to provide fault tolerance and load balancing for the Active Directory data store. The other computers on the network normally run a client operating system, such as Windows 2000 Professional, Windows NT Workstation, Windows 95, or Windows 98 (although they may utilize Windows 2000 Server or other operating systems). As in Windows NT 4.0 domains, resources, such as hard disks and printers, can be shared from any computer on the network.

What Is Active Directory?

Active Directory is the directory service used by Windows 2000. It is a core new feature of the Windows 2000 operating systems.

A *directory service* consists of two parts — a centralized, hierarchical database that contains information about users and resources on a network, and a service that manages the database and enables users of computers on the network to access the database. In Windows 2000, the database is called the Active Directory data store, or sometimes just the directory. The Active Directory data store contains information about various types of network objects, including printers, shared folders, user accounts, groups, and computers. Windows 2000 Server computers that have a copy of the Active Directory data store, and that run Active Directory are called *domain controllers*. In a Windows 2000 domain, a read/write copy of the Active Directory data store is physically located on each domain controller in the domain. A *domain* is a logical grouping of networked computers in which one or more of the computers has shared resources, such as a shared folder or printer, and in which all of the computers share a common Active Directory data store.

The three primary purposes of Active Directory are:

- To provide user logon and authentication services
- To enable administrators to organize and manage user accounts, groups, and network resources
- To enable authorized users to easily locate network resources, regardless of where they are located on the network

So why is Active Directory so cool? I'll answer that question in the next section by discussing some of the features of Active Directory.

4. What is Active Directory? Write the primary purposes of Active Directory.

Flexibility and Extensibility

Active Directory can evolve as your business does. It is not a static structure that, once implemented, can never be changed.

Active Directory is said to be *extensible*. This means that new classes of objects can be added, and new attributes can be added to classes of objects already present.

Now that you have a basic understanding of what Active Directory is and an awareness of some of its key benefits, it's time to wade in a little deeper to the actual structure of Active Directory and its many components.

Understanding the Structure of Active Directory

To review: Active Directory has a hierarchical, tree-like structure. Information about network users and resources is stored in the *Active Directory data store*, which is a structured, centralized database. A read/write copy of the Active Directory data store is physically located on each domain controller in a Windows 2000 domain. This data store is commonly referred to as the directory.

In order to talk in greater depth about the structure of Active Directory, I need to introduce and define several new terms. Many of these terms are components of Active Directory, and some of the terms are used to define relationships between the components. In the following sections I'll discuss objects and classes, schema, the global catalog, and the hierarchical structure of Active Directory, including domains, organizational units, trees, trust relationships, and forests. I'll also discuss Active Directory names and naming conventions, as well as security. When I'm finished, you'll have a much better picture of how Active Directory is structured.

5. Elaborate the Structure of Active Directory.
Objects and Classes, Schema, Global Catalog, Domains, Organizational Units, Trees

Objects and Classes

An Active Directory *object* is a record in the directory that is defined by a distinct set of attributes. The attributes of an object are the same as the object's properties. The terms are synonymous; however, the term *properties* is more prevalent throughout the Windows 2000 user interface.

The specific attributes that an object can have are defined by the object's class. A *class* is simply a template that is used to define the attributes of an object when it is created. A class defines the required and optional attributes of the objects that are instances of that class. For example, the Computer class contains a list of the required and optional attributes that are used when a computer object is created. All computer objects will be created using the same Computer class definition.

There are many classes of Active Directory objects. Some of the classes are:

- Computer
- Contact
- Group
- Organizational Unit
- Domain
- Printer
- User
- Shared Folder

Schema

In Active Directory terminology, the *schema* is a formal definition — a set of rules, if you wish — of all of the classes of objects and their attributes that are stored in the directory. The schema governs the structure of the directory, including how various objects in the directory fit into the directory's hierarchical structure.

The schema is what makes Active Directory extensible. As organizations change, it may be necessary to add or modify object attributes, or even to create new classes. The use of certain applications, in particular, may require these kinds of modifications. Microsoft anticipates that application vendors will provide the means to modify the schema when necessary to support their application's specific requirements.

Windows 2000 Server includes a tool to modify the schema. It is a Microsoft Management Console (MMC) snap-in that is only available after installing the Windows 2000 Administration Tools (ADMINPAK) on a Windows 2000 computer. The name of the snap-in is Active Directory Schema.

Because the schema is the heart of Active Directory, it's important that it be protected from accidental or unauthorized modification. For this reason, Microsoft created a special Security Group for Windows 2000 called Schema Admins. Only users with this permission can run programs that will modify the schema.

Global Catalog

The *global catalog* is a master, searchable index that contains information about every object in every domain in a forest. For now, you can think of a *forest* as all of the domains that make up a company's network. Forests will be covered in more technical detail later in this chapter.

The global catalog, in conjunction with various search tools, is what enables administrators and users to search for and quickly locate an object, regardless of where the object is located on the network.

Windows 2000 automatically creates, by default, a global catalog on the first domain controller that is installed in a forest. You can configure other domain controllers to maintain a copy of the global catalog, as well. The global catalog contains a full copy, or replica, of all objects in its host domain, and a partial replica of all objects in all other domains in the forest. A partial replica includes the most common properties of every object, but not all of the properties of every object.

Hierarchical Structure

By now you've read the term "hierarchical structure" a zillion times. But what does it mean, exactly? A *hierarchical structure* refers to a manner of organizing a group of interrelated elements in which the elements are ranked or stacked, one above the other. An example of a hierarchical

structure that you are probably familiar with is an organizational chart. Figure 2-1 shows an organizational chart for ABC Bank.

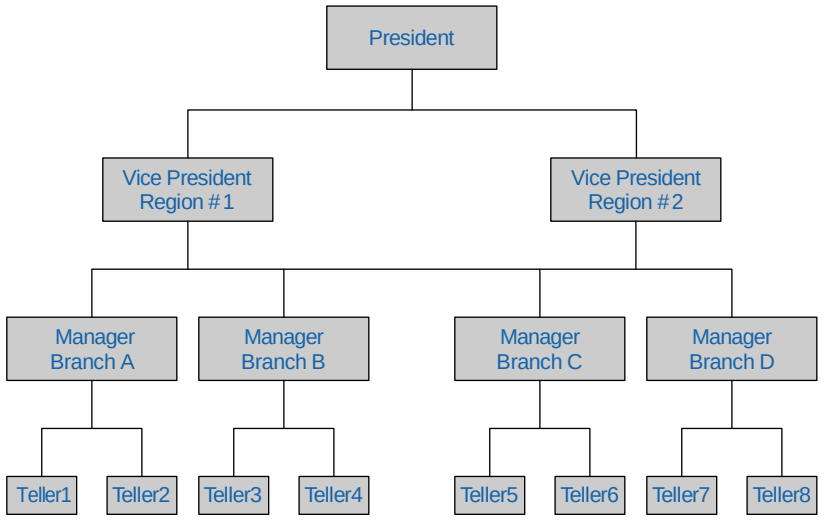

FIGURE 2-1 Organization chart of ABC Bank

In the organizational chart of this regional bank, the President is at the top of the chart, and beneath the President is a level consisting of two Vice Presidents. After the Vice President level is a layer of management staff, and beneath this layer is a level that represents the numerous bank tellers. The hierarchical structure typically has at the top a single element, which branches into lower layers that contain progressively more elements the farther down you go.

The key building blocks in the Active Directory hierarchical structure are domains, which are the focus of the next section.

Domains

Domains are the fundamental units that make up Active Directory. As stated previously, a *domain* is a logical grouping of networked computers in which one or more of the computers has shared resources, such as a shared folder or printer, and in which all of the computers share a common Active Directory data store that contains user account, resource, security, and other information. Active Directory consists of one or more domains.

A domain is a natural security boundary in a Windows 2000 network. Users from other domains cannot pierce this boundary to access shared resources unless trust relationships are created between the domains to

permit user access. More information about trust relationships is provided later in this chapter.

A domain can span several geographic locations of a company, or a domain can be created for every location. Sometimes the needs of the departments, divisions, or subsidiaries of an organization determine the number and structure of the domains needed to effectively manage the organization's network.

The domains that make up Active Directory usually correspond to the network's DNS domains, and typically use the same FQDN naming convention used by DNS servers. FQDN stands for *fully qualified domain name*, and is the naming convention used on the Internet. The format for an FQDN is *server_name.domain_name.root_domain_name*. I'll discuss names and naming conventions in a bit more detail later in this chapter.

Domains contain objects, and can also contain organizational units, which are discussed in the next section.

Organizational Units

Organizational units are a type of Active Directory object, and are sometimes called container objects. They contain objects and other organizational units from their own domain. Organizational units are often called by their abbreviated name (OUs.

An organizational unit is used to organize related objects and other organizational units in Active Directory in much the same way that a folder is used to organize related files and other folders in a volume. Also, the organizational unit is the smallest container component of Active Directory to which you can delegate administrative authority or assign group policy. The primary purpose of an organizational unit, then, is the organization of related objects and other organizational units to simplify administration.

For example, suppose an administrator wants to delegate network administration of the Sales department to an assistant administrator. The administrator decides to group together all of the objects associated with the Sales department (including users, computers, printers, shared folders, and groups). Then the administrator creates an organizational unit and

places all of the objects associated with the Sales department into this organizational unit. Completing these steps enables the administrator to delegate administration for the Sales department by assigning the assistant administrator the permissions required to administer the organizational unit and its contents.

Trees

In Active Directory terminology, a *domain tree* is a hierarchical grouping of one or more domains that must have a single root domain, and may have one or more child domains. In a domain tree, the root domain is the domain at the top (or root) of the tree.

Domains in a domain tree are often spoken of in terms of parent domains and child domains. A *parent domain* is any domain that is above another domain in the domain tree hierarchy. A *child domain* is any domain that is below another domain in the tree. A domain can be a parent to a domain below it and a child to the domain above it. In a multidomain tree, the root domain is always a parent domain. Figure 2-2 illustrates a domain tree. Notice that there is only one root domain in the tree, but that the tree contains more than one child domain.

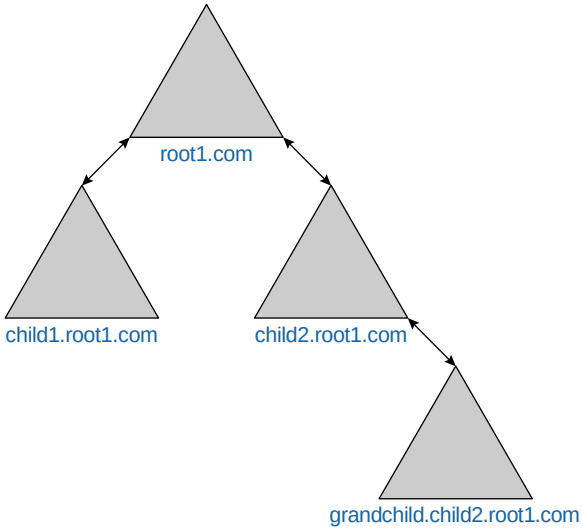

FIGURE 2-2 A domain tree

Also notice the naming structure used in Figure 2-2. In a domain tree, the domains that make up the tree have contiguous DNS domain names. The root domain's name forms the basis of (and will be a part of) the

6. Define Trust Relationship. Describe different types of Trust Relationships in Windows 2000.

FQDNs of all of the other domains in the tree. A child domain's FQDN is created by appending the name of the parent domain to its own NetBIOS name by using the *child_domain.parent_domain.root_domain.* `com` format. For example, in Figure 2-2, the domain with a NetBIOS name of child1 appends the name of its parent domain, `root1.com`., to its own name, resulting in an FQDN of `child1.root1.com`. The root domain in a domain tree also takes its name in this way, by appending the name of the first-level DNS domain that it is a member of to its own NetBIOS name. In Figure 2-2, the root domain with a NetBIOS name of root1 appends the name of the first-level DNS domain, `com`, to its own name, resulting in an FQDN of `root1.com`.

In organizations that require multiple domains, a domain tree enables any permitted user in any domain in the tree to access shared resources in any domain in the tree.This user access is made possible by the special trust relationships that exist between the domains in the tree.

Trust Relationships

To manage the interaction between multiple domains, trust relationships are necessary. A *trust relationship*, or *trust*, is an agreement between two domains that enables users in one domain to be authenticated by a domain controller in another domain, and therefore to access shared resources in the other domain.

The terminology used to discuss trusts is sometimes confusing, so a good portion of this section is dedicated to explaining and clarifying these terms. Once you've mastered the terminology, trust concepts are much easier to understand.

Trusting Domain vs. Trusted Domain Two terms are commonly used to refer to a trust between two domains: trusting domain and trusted domain. The *trusting domain* is the domain that has resources to share with user accounts in the trusted domain. The trusting domain trusts the trusted domain. The *trusted domain* is the domain that contains the user accounts that want to access the shared resources in the trusting domain.The trusted domain is trusted by the trusting domain.

A trust relationship between two domains is depicted in diagrams by using an arrow to point from the trusting (resource) domain to the trusted (user accounts) domain. Figure 2-3 illustrates a trust relationship between the `west.com` domain and the `east.com` domain.The `west.com` domain

is the trusting domain, and the east.com domain is the trusted domain. Notice that the arrow points toward the domain with the user accounts.

This trust relationship enables users from the east.com domain to access shared resources located in the west.com domain.

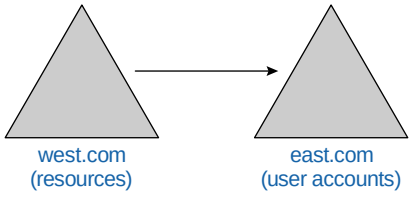

FIGURE 2-3 The west.com domain trusts the east.com domain

Intransitive and Transitive Trusts An *intransitive trust* is a trust relationship between two domains that does not extend beyond these two domains to other domains. An intransitive trust is a *one-way trust*, meaning that a single trust relationship exists between the two domains.

Suppose that the a.com domain trusts the b.com domain. Further suppose that the b.com domain trusts the c.com domain. Figure 2-4 shows these trust relationships.

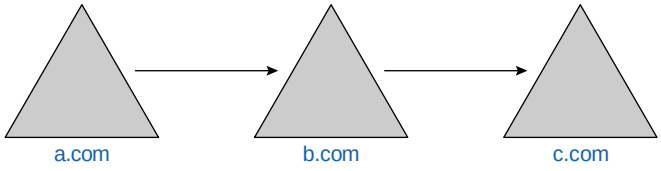

FIGURE 2-4 Intransitive trusts

At first glance, it might appear that the user accounts in the c.com domain are able to access resources in the a.com domain, but this is not the case. A trust relationship does not exist between the a.com domain and the c.com domain. Therefore, users in the c.com domain can't access resources in the a.com domain.

It is possible to establish a two-way trust relationship between two domains by creating two, one-way trusts between those domains. In a *two-way trust relationship*, two domains trust each other.

A *transitive trust* is a trust relationship between two Windows 2000 domains in the same domain tree (or forest) that can extend beyond these two domains to other trusted domains within the same domain tree (or forest).A transitive trust is always a two-way trust, meaning that both of

the domains trust each other. By default, all Windows 2000 trusts within a domain tree (or forest) are transitive trusts.

Transitive trusts are depicted in diagrams by a single line with an arrow at each end. Figure 2-5 illustrates transitive trusts in a Windows 2000 domain tree.

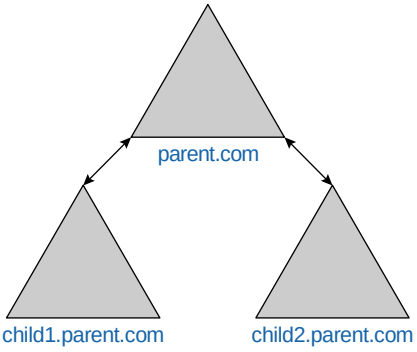

FIGURE 2-5 Transitive trusts

Notice that in Figure 2-5, transitive trust relationships exist between each child domain and the parent domain, but that no trust relationship exists directly between the two child domains. Nonetheless, the transitive trust relationships make it possible for users in the `child1.parent.com` domain to access resources located in both the `parent.com` domain and in the `child2.parent.com` domain. Likewise, users in the `child2.parent.com` domain can access resources located in both the `parent.com` domain and in the `child1.parent.com` domain because of the transitive trusts that connect the three domains.

In Windows 2000 domain trees, Windows 2000 Server automatically creates two-way, transitive trust relationships between a parent domain and a child domain when the child domain is created in the domain tree.The presence of transitive trust relationships between all of the domains in a Windows 2000 domain tree makes it possible for a user in one domain to access a shared resource located in any domain in the tree, regardless of how many domains separate the user and the shared resource.

Windows NT Server 4.0 doesn't support transitive trusts — it only supports intransitive trusts. This means that the only type of trust relationship possible between a Windows 2000 domain and a Windows NT domain is an intransitive trust.

Explicit Trusts An *explicit trust* is a trust that an administrator creates, versus a trust that is automatically created by Windows 2000. An explicit trust can be either transitive or intransitive. Explicit trusts are sometimes used when you need to manage trusts between a Windows 2000 domain and a Windows NT domain. Explicit trusts are also used in large, multidomain forests to shorten the path between two domains to shorten the time required for authentication and logon.

Forests

Earlier in this chapter I said you could think of a forest as being all of the domains that compose a company's network. A more technically accurate definition of a forest is a group of one or more domain trees, linked by transitive trusts, that shares a common schema and global catalog.

A forest begins with one domain and one domain tree. It's kind of a difficult concept to grasp, but when you install Active Directory on the first domain controller on your network, Windows 2000 creates a domain, a domain tree, and a forest all at the same time. So, even though you've only installed Active Directory on one computer, you've got all of these big-picture elements created and ready to go. Now the forest can grow as you add additional domains and domain trees.

Figure 2-6 illustrates a forest that consists of two domain trees. Notice that this forest contains two root domains, each of which forms the basis for its own domain tree. Also notice that a single, transitive trust connects the two domain trees.

Take another look at Figure 2-6, and notice the domain names. By definition, the domains in a domain tree have contiguous DNS domain names. In this example, `rootA.com` is contained in the name of every domain in its tree. Likewise, `rootB.org` is contained in the name of every domain in its tree. However, that's as far as it goes. The two domain trees themselves do not have contiguous DNS domain names, even though they have been joined together in a forest.

7. Define GUID, SID, RDN, DN, UPN

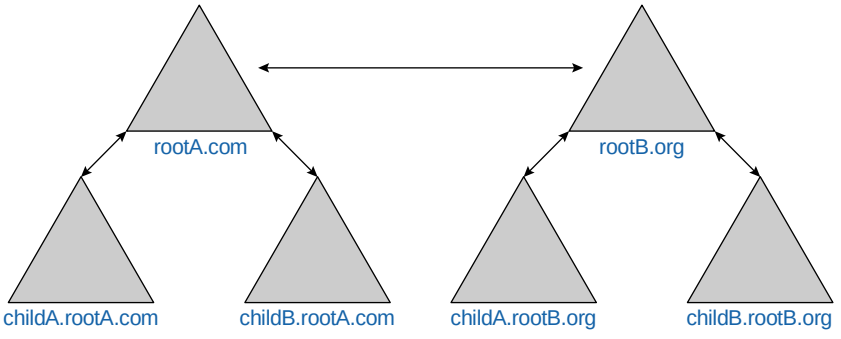

FIGURE 2-6 A forest

A forest takes its name from the root tree, which is the first tree created in the forest.

Names and Naming Conventions

Names are of critical importance in Active Directory. In this section, I'll explain the types of names and naming conventions used by Active Directory.

Within Active Directory, each object has a name. When you create an object in Active Directory, such as a user or a computer, you assign the object a name.This name must be unique within the domain — you can't assign an object the same name as any other object (regardless of its type) in that domain. If you have a user named AlanC, for example, you can't create a computer account in the domain that is also named AlanC.

For more information on developing names for domains, organizational units, users, groups, and computers, see the section titled "Planning Naming Conventions" later in this chapter.

At the same time that you create an object, not only do you assign a name to the object, but Active Directory also assigns identifiers to the object. Active Directory assigns every object a globally unique identifier (GUID), and assigns many objects a security identifier (SID). A *GUID* is typically a 32-digit hexadecimal number that uniquely identifies an object within Active Directory. A *SID* is a unique number created by the Windows 2000 Security subsystem that is assigned only to *security principal objects* (users, groups, and computers) when they are created. Windows 2000 uses SIDs to grant or deny a security principal object access to other objects and network resources.

Active Directory uses a hierarchical naming convention that is based on Lightweight Directory Access Protocol (LDAP) and DNS standards.

8. Objects in the Active Directory are referenced by Active Directory name types. Explain all the types with proper example.

Objects in Active Directory can be referenced by using one of three Active Directory name types:

- Relative distinguished name (RDN)
- Distinguished name (DN)
- User principal name (UPN)

A *relative distinguished name* (RDN) is the name that is assigned to the object by the administrator when the object is created. For example, when I create a user named AlanC, the RDN of that user is AlanC. The RDN only identifies an object — it doesn't identify the object's location within Active Directory. The RDN is the simplest of the three Active Directory name types, and is sometimes called the common name of the object.

A *distinguished name* (DN) consists of an object's RDN, plus the object's location in Active Directory. The DN supplies the complete path to the object. An object's DN includes its RDN, the name of the organizational unit(s) that contains the object (if any), and the FQDN of the domain. For example, suppose that I create a user named AlanC in an organizational unit called US in a domain named `Exportsinc.com`.The DN of this user would be:

```
AlanC@US.Exportsinc.com
```

A *user principal name* (UPN) is a shortened version of the DN that is typically used for logon and e-mail purposes.A UPN consists of the RDN plus the FQDN of the domain. Using my previous example, the UPN for the user named AlanC would be:

```
AlanC@Exportsinc.com
```

Another way you can think of a UPN is as a DN stripped of all organizational unit references.

Security

As mentioned previously, Active Directory resides in the Windows 2000 Security subsystem. Together, Active Directory and the Security subsystem protect Active Directory against unauthorized access. The Active Directory/Security subsystem team uses access control lists (ACLs) to determine who can access (and/or modify) an object. An ACL is a list of SIDs and the associated access privileges assigned to each SID. Each object and network resource has an ACL associated with it.

Working with File Systems

Before you attempt to configure a computer's disks, it's important that you have a clear understanding of the different file systems that Windows 2000 supports. Windows 2000 supports five file systems: the file allocation table (FAT) file system, the FAT32 file system, the Windows NT file system (NTFS), the Compact Disc File System (CDFS), and the Universal Disk Format (UDF). Table 6-1 shows which file systems are supported by various operating systems.

TABLE 6-1 File System Support by Operating System

Operating System	File Systems Supported
Windows 2000	FAT, FAT32, NTFS, CDFS, UDF
Windows NT 4.0	FAT, NTFS, CDFS, UDF
Windows NT 3.51 (and earlier versions)	FAT, NTFS, CDFS, HPFS
Windows 98	FAT, FAT32, CDFS, UDF
Windows 95	FAT, (FAT32 on OSR2 only), CDFS, UDF
Windows 3.*x* and 3.1*x*	FAT, CDFS
OS/2 1.*x*	FAT, CDFS, HPFS
MS-DOS	FAT, CDFS

In the following sections I'll describe each of the file systems supported by Windows 2000 in detail. I'll also discuss the capabilities and limitations of each of these file systems.

FAT

The *file allocation table (FAT) file system* used by Windows 2000 is a modified version of the FAT file system used by MS-DOS. FAT (sometimes called FAT16) is the only hard disk file system supported by Windows 95 (versions prior to OSR2), Windows 3.*x*, Windows 3.1*x*, and MS-DOS. So, if you want to configure a Windows 2000 computer to dual boot between Windows 2000 and Windows 95 (versions prior to OSR2), Windows 3.1*x*, or MS-DOS, your computer's first partition on the first hard disk must use the FAT file system.

9. Describe FAT file system at length.
Basic Description, Security, Naming Conventions, Speed and Access to Files, Volume Size

If you're not sure whether you have an OSR2 version of Windows 95, there's an easy way to find out. From the Windows 95 desktop, select Start ➪ Settings ➪ Control Panel. Then double-click the Systems application, and examine the General tab, which lists specific information about the system installed on your computer. If your system version is 4.00.950 or 4.00.950 a, then you have a version of Windows 95 that was released prior to OSR2.

Now I'll give a brief overview of the characteristics and features of the FAT file system, including security, naming conventions, speed of access to files, and volume size.

Security

The FAT file system does not support file and folder security in Windows 2000. Because file and folder security is not supported on a FAT volume, any user who is logged on locally to a computer has full control of all of the files and folders located in the FAT volume(s) on that computer. This applies only to local access.

However, you can use share permissions to control users' access to shared folders over the network. Share permissions affect only the access of files and folders over the network, not when someone is logged on locally. So, if you need local file and folder security, you should use the NTFS file system instead of the FAT file system.

Naming Conventions

The FAT file system, as used by Windows 2000, supports the use of long filenames. Filenames can be up to 255 characters in length.

Filenames can contain any character except \ / : * ? " < > | and may begin with any permitted character. Filenames can contain spaces and multiple periods, and the characters after the last period are considered the filename extension.

10. Describe FAT32 file system at length.
Basic Description, Security, Naming Conventions, Speed and Access to Files, Volume Size

The FAT file system preserves uppercase and lowercase in filenames, but filenames are not case sensitive. Because of this, I can request the file `ALAN.DOC` by typing **Alan.doc**, **ALAN.DOC**, or **alan.doc**, and Windows 2000 always retrieves `ALAN.DOC`.

Speed of Access to Files

Access speed to files in a FAT volume is dependent on many factors, includ- ing volume size, number of files in a folder, and fragmentation.

Windows 2000 accesses files in FAT volumes smaller than 512MB faster than it accesses files in similar-sized FAT32 and NTFS volumes.

If the volume size is larger than 512MB, however, or when there is a large number of files in a folder, Windows 2000 accesses files in FAT32 and NTFS volumes much faster than it accesses files in a FAT volume of similar size.

Windows 2000 usually accesses files in a highly fragmented FAT volume more slowly than it accesses files in an NTFS volume of similar size.

Volume Size

The maximum size of a FAT volume on all operating systems except Windows 2000 and Windows NT is 2GB. Both Windows 2000 and Win- dows NT support FAT volumes up to 4GB. This is possible because Windows 2000 and Windows NT support a larger cluster size (up to 64K) than do other operating systems.

The maximum size of a file in a FAT volume is 4GB.The FAT file system, as used by Windows 2000, does not support file compression.

FAT32

The *FAT32 file system* used by Windows 2000 is the same as the FAT32 file system that was released with Windows 95 OSR2 and Windows 98. The FAT32 file system is only supported by Windows 2000, Windows 98, and Windows 95 OSR2.

If you want to dual boot between Windows 2000 and Windows 98 (or Windows 95 OSR2), you can use either the FAT32 or FAT file system on your computer's first volume.

In the sections that follow I'll cover the specific characteristics of the FAT 32 file system, including security, naming conventions, speed of access to files, and volume size.

Security

Like the FAT file system, the FAT32 file system does not support file and folder security in Windows 2000. Because file and folder security is not supported on a FAT32 volume, any user who is logged on locally to a computer has full control of all of the files and folders located in the FAT32 volume(s) on that computer. This applies only to local access.

However, you can use share permissions to control users' access to shared folders over the network. Share permissions affect only the access of files and folders over the network, not when someone is logged on locally. So, if you need local file and folder security, you should use the NTFS file system instead of the FAT32 or FAT file systems.

Naming Conventions

The naming conventions supported by the FAT32 file system are identical to those supported by the FAT file system:

- Filenames can be up to 255 characters in length.
- Filenames can contain any character except \ / : * ? " < > | and may begin with any permitted character. Filenames can contain spaces and multiple periods.
- The FAT32 file system preserves uppercase and lowercase in filenames, but filenames are not case sensitive.

Speed of Access to Files

Access speed to files in a FAT32 volume is dependent primarily on volume size and fragmentation.

Windows 2000 accesses files in FAT32 volumes larger than 512MB faster than it accesses files in similar-sized FAT volumes, but slower than it accesses files in similar-sized NTFS volumes.

Windows 2000 usually accesses files in a highly fragmented FAT32 volume more slowly than it accesses files in an NTFS volume of similar size.

Volume Size

Although the maximum size of a FAT32 volume on Windows 98 and Windows 95 OSR2 is 2 terabytes (TB), the disk management utilities contained in Windows 2000 only enable you to create and format a FAT32 volume up to 32GB. Windows 2000 does support FAT32 volumes larger than 32GB that are created by other operating systems.

The maximum size of a file in a FAT32 volume is 32GB. Like the FAT file system, FAT32 does not support file compression.

11. Describe NTFS file system at length.
Basic Description, Security, Naming Conventions, Speed and Access to Files, Volume Size

NTFS

The *Windows NT file system (NTFS)* is the most powerful file system supported by Windows 2000. Only Windows 2000 and Windows NT support NTFS — no other Microsoft operating systems currently support this file system.

Windows 2000 NTFS is a newer version than Windows NT NTFS, and supports several features not supported by Windows NT NTFS. Because of this, if you want to dual boot between Windows 2000 and Windows NT, you must have Windows NT 4.0 with Service Pack 4 or later installed.

When it comes to security, naming conventions, speed of access to files, and volume size, NTFS in Windows 2000 has its own unique characteristics. Additionally, NTFS has some features not supported by the FAT or FAT32 file systems.

Security

NTFS provides file and folder security for both local and remote users on a network. NTFS is the only file system discussed here that permits the assigning of permissions to individual files and folders.

So how does NTFS security actually work? NTFS security controls access to files on an NTFS volume by utilizing the user's security identifier (SID) to determine which files that user can access. Each file and folder on an NTFS volume has an access control list (ACL) associated with it. The *ACL* is a list that contains user and group SIDs, with the associated privileges of each user and group.

NTFS supports the *Encrypting File System* (EFS). EFS enables you to store files on an NTFS partition in an encrypted format so that even if an unauthorized user removes a hard drive from your computer, that user will be unable to access the sensitive data contained in the encrypted file.

In addition to the security provided by NTFS, remember that because Windows 2000 requires a user to log on before accessing files, Windows 2000's security is greater than operating systems that don't require the user to log on.

Naming Conventions

Like the FAT and FAT32 file systems, NTFS supports the use of long filenames. Filenames can be up to 255 characters in length.

Filenames can contain any character except `\ / : * ? " < > |` and may begin with any permitted character. Filenames can contain spaces and multiple periods, and the characters after the last period are considered the

filename extension.

NTFS preserves uppercase and lowercase in filenames. Filenames are not case sensitive (except when used by a POSIX application). For example, a Win32 application does not distinguish between `Money.DOC`, `MONEY.DOC`, and `money.doc` — it treats all three names as though they were the same file.

The POSIX subsystem, however, is case sensitive with respect to filenames, because it does not translate a request for a file into all uppercase letters as the Win32 and other subsystems do. A POSIX application treats the filenames in the previous paragraph as though they were three separate files: `Money.DOC`, `MONEY.DOC`, and `money.doc`. You must use a POSIX application if you want to access these three different files — if you attempt to access `Money.DOC` with a Win32 application (no matter how you type the file name), you will always retrieve the `MONEY.DOC` file because the Win32 Subsystem translates file requests into all uppercase letters.

Speed of Access to Files

NTFS usually provides faster access than the FAT or FAT32 file systems to files stored on a large volume that contains many files. NTFS is able to access files in this situation faster than the FAT or FAT32 file systems because NTFS uses an enhanced binary tree to locate files. A binary tree search is a faster mechanism for searching through a large number of filenames than the sequential read mechanism used on FAT and FAT32 volumes.

Volume Size

The maximum theoretical size of an NTFS volume is 16 exabytes (an *exabyte* is one billion billion bytes, or a giga-gigabyte). However, when you actually implement NTFS on current standard industry hardware, there is a functional limitation of 2TB.

The maximum size of a file in an NTFS volume is limited only by the amount of free space in the NTFS volume.

Additional Features Not Supported by FAT or FAT32

NTFS has several other unique attributes and features that are not found in, nor supported by, the FAT or FAT32 file systems.

- NTFS supports a compression attribute for each file. You can choose which files to compress and which ones to leave uncompressed. The compression algorithm NTFS uses is similar to the one used by Drivespace in MS-DOS. Using compression provides an approximately 40 to 50 percent increase in hard disk space.

- NTFS is a highly reliable, recoverable file system. It is not necessary to periodically run `Chkdsk.exe` on an NTFS volume.

- Using NTFS greatly reduces fragmentation on volumes. However, files can still become fragmented when their size is increased. Windows 2000 (unlike Windows NT) includes a defragmentation utility which can be used to defragment FAT, FAT32, and NTFS volumes.
- NTFS maintains a recycle bin for each user.
- NTFS enables you to mount a volume on a folder in a different volume. The term *mounting a volume* refers to a disk management technique sometimes used to access space on more than one hard disk (or volume) but still retain and use a single drive letter. The result of this feature is that a folder's contents are physically stored on a different hard disk (or volume), but this folder and its contents appear to users to be located in the current volume. This feature produces results similar to those produced by executing the `mount` command on a UNIX computer.
- NTFS supports the Encrypting File System (EFS).
- NTFS supports disk quotas. *Disk quotas* is a volume management tool that is enabled on a volume-by-volume basis. Once enabled, disk quotas automatically track disk space usage on a user-by-user basis, and prevent individual users from exceeding the disk space limitations that they have been assigned by administrators.

The first four features in the preceding list are supported by both Windows 2000 NTFS and Windows NT NTFS. The last three features are new features that are supported only by Windows 2000 NTFS.

A couple of final tidbits about NTFS:

- You can't use NTFS to format floppy disks.
- You can change media in a removable media device (such as a Zip drive) that has been formatted with NTFS without rebooting the computer. (This feature was not supported by Windows NT.)

Which File System Should I Use?

Because of its speed, security, and recoverability, I recommend the use of NTFS on all volumes except for floppy disks, and volumes that are used to dual boot between Windows 2000 and another operating system.

12. Briefly describe CDFS, UDF and HPFS file systems supported by Windows 2000.

If you require dual boot, and the other operating system supports FAT32, then I recommend FAT32 over FAT because of FAT32's speed and support of larger volume sizes.

CDFS

The *Compact Disc File System (CDFS)* supports access to compact discs. It is not used on a computer's hard disks — this file system is used only on CD-ROM devices that read and/or write compact discs. Because of the prevalence of CD-ROM devices, CDFS is supported by most operating systems.

UDF

The *Universal Disk Format (UDF)* is a file system used to access read-only digital video discs (DVDs). Like CDFS, this file system is not used on a computer's hard disks — this file system is used only on DVD-ROM devices.

HPFS

Windows 2000 does not support the high performance file system (HPFS), although some of the earliest versions of Windows NT did. If you want to upgrade to Windows 2000 from an early version of Windows NT that used HPFS, you must convert your HPFS volume to NTFS before performing the upgrade.

Converting from FAT or FAT32 to NTFS

In Windows 2000 you can format a new volume with either FAT, FAT32, or NTFS. But what do you do when you want to change the file system on an existing volume? You can change an existing FAT or FAT32 volume into an NTFS volume by using `Convert.exe`. This is a fairly simple procedure. When you use `Convert.exe` all data on the existing volume is retained..

However, it is a one-way process — there is no way to convert an NTFS volume into a FAT or FAT32 volume without first backing up, reformatting the volume, and restoring the data.

To convert a FAT or FAT32 volume into an NTFS volume, use the `Convert.exe` command at a command prompt. To start a command prompt, select Start ➪ Programs ➪ Accessories ➪ Command Prompt. The syntax for the `Convert.exe` command is:

```
CONVERT volume /FS:NTFS [/V]
```

The following is an explanation of this syntax:

- *Volume* This specifies the drive letter (followed by a colon) or mount point to convert to NTFS.
- `/FS:NTFS` This indicates that the file system should be converted to NTFS. This is an outdated switch, because NTFS is the only file system that you can use `Convert.exe` to switch to in Windows 2000; but its use, in terms of command syntax, is still required.
- `/V` This optional switch specifies that `Convert.exe` will run in *verbose mode*. Running a command in verbose mode will display the maximum amount of information and detail to the user.

Let me illustrate the use of this command with a couple of examples:

1. To convert drive `D:` from FAT to NTFS, use the following command:

   ```
   CONVERT D: /FS:NTFS
   ```

2. To convert a mount point named `C:\Data` from FAT32 to NTFS, using the optional verbose mode, use the following command:

   ```
   CONVERT C:\Data /FS:NTFS /V
   ```

To successfully use the `Convert.exe` command, `Convert.exe` must be the *only* application that accesses the drive or mount point you want to change during the conversion process. If Windows Explorer accesses the drive or mount point you are trying to convert, if you are trying to convert the boot partition, or if your active command prompt has the drive you are trying to convert as its current drive, Windows 2000 will display an error message stating that `Convert.exe` cannot gain exclusive access to the

drive or mount point, and asks if you want to schedule it to be converted the next time the system restarts.

If you try to execute the `Convert.exe` command, but can't gain exclusive access to a drive or mount point, type **Y** when asked if you want to schedule it to be converted the next time the system restarts. Windows 2000 will convert the file system when you restart your computer.

Creating and Managing User Accounts

User accounts are records that contain unique user information, such as user name, password, and any logon restrictions. User accounts enable users to log on to Windows 2000 computers, and to access resources on the network.

Built-in User Accounts

There are two Windows 2000 built-in user accounts: Administrator and Guest. On nondomain controllers, the built-in user accounts are created automatically during the installation of Windows 2000. On a domain controller, the built-in user accounts are created automatically during the installation of Active Directory.

The Administrator user account has all of the rights and permissions needed to fully administer a Windows 2000 computer or a Windows 2000 domain. The Administrator account can be used to perform numerous tasks, such as creating and managing users and groups, managing file and folder permissions, and installing and managing printers and printer security.

The Administrator account, because of its powerful capabilities, can pose a security risk to your network if a nonauthorized user is able to guess the password for the account. For this reason, you should consider renaming the Administrator account. (I'll explain how to rename a user account later in this chapter.)

You can't delete the Administrator account. You also can't disable the Administrator account, nor can you remove this account from the Administrators local group. Incidentally, it's the Administrator account's membership in the Administrators local group that gives the Administrator account all of its rights and permissions.

The Guest account, which is disabled by default, is designed to permit limited access to network resources to occasional users who don't have their own user accounts. For example, a client visiting your office might

want to connect a laptop computer to your network in order to print a document. Once the Guest account is enabled, the client can log on using this account. You can specify, in advance, which network resources are available to the Guest account by assigning the appropriate file, folder, and printer permissions to this account.

The Guest account does not require a password. If your network contains sensitive data, I recommend, for security reasons, that you leave the Guest account disabled. In this situation, instead of using the Guest account, you should establish a user account for each and every person who needs access to network resources.

You can't delete the Guest account, but you can rename it.

Creating User Accounts

Every person who uses the network on a regular basis should have a user account.

There are two kinds of user accounts: local user accounts and domain user accounts. *Local user accounts* enable users to log on to the local computer and to access that computer's resources. *Domain user accounts* enable users to log on to the domain and to access resources in the domain.

In order to create local user accounts, you must be a member of either the Administrators or Power Users group on the local computer. In order to create domain user accounts, you must be a member of either the Administrators or Account Operators group in the domain.

I'll show you how to create user accounts in just a minute, but before I do, I want to say a few words about naming conventions and passwords.

Naming Conventions

When you create user accounts, keep in mind a few simple rules for user names:

- User names (which are referred to as *user logon names* in Active Directory Users and Computers) can be from one to 20 characters long.

- User names must be unique. A domain user name can't be the same as another user, group, or computer name within the domain. A local user name can't be the same as another user, group, or computer name within the local computer's account database.
- The following characters may *not* be used in user names:
  " / \ [ ] : ; | = , + * ? < >
  In addition, a user name can't consist entirely of spaces or periods.

If you have more than a few people in your organization, it's a good idea to plan your user account naming convention.

There are probably as many user account naming schemes as there are network administrators. Sometimes the overall length of a user name is limited to eight characters, so that the name is compatible with MS-DOS directory name limitations. While this eight-character limitation is com- mon, it's certainly not mandatory, especially on most of today's networks. A few common naming conventions for user names include:

A. The first seven letters of the user's first name plus the first letter of the user's last name
B. The first letter of the user's first name plus the first seven letters of the user's last name
C. The user's initials plus the last four digits of the user's employee number
D. Various hybrid combinations of the preceding schemes

Table 9-1 shows how three user names would appear using the naming conventions described in A, B, and C.

TABLE 9-1 Common User Account Naming Conventions

Full Name	Scheme A	Scheme B	Scheme C
Nadine Smith	NadineS	NSmith	NS5500
Robert Jones	RobertJ	RJones	RJ1234
Jonathan Whitmore	JonathaW	JWhitmor	JW2266

In addition to choosing a naming convention, you should have a way to handle exceptions. It's quite common, for example, for two users to have the same first name and last initial, such as Mike Smith and Mike Sutherland. If your company uses the naming convention described in scheme A, you would need to resolve the potentially duplicate user names

for these two employees. You could resolve the problem by assigning Mike Smith the user name of MikeS (assuming he was hired before Mike Sutherland), and assigning Mike Sutherland the user name of MikeSu.

Passwords

I'll just say a few words about passwords. Everyone knows that using passwords protects the security of the network, because only authorized users can log on.

When user accounts are created, you should have a plan for managing passwords. Will passwords be assigned and maintained by the network administrator? Or, will users choose their own passwords?

When users maintain their own passwords, it's a good idea to remind them of a few password security basics:

- Don't use your own name or the name of a family member or pet as a password. (This is a common security loophole in most networks.)
- Never tell your password to anyone.
- Don't write your password on a sticky note and then stick it on your monitor. Other not-so-hot places to store your password are on or under your keyboard; in your top desk drawer; in your Rolodex; or in your briefcase, wallet, or purse.
- Use a sufficiently long password. I recommend using eight or more characters in a password. The longer the password, the more difficult it is to guess.
- Use a mix of uppercase and lowercase letters, numbers, and special characters. Remember, passwords are case-sensitive.
- If passwords are required to be changed regularly, don't use the same password with an incremental number at the end, such as Alan01, Alan02, Alan03, and so on. (Don't laugh. This may seem like common sense, but I've seen several network administrators actually do this.)

Creating and Managing Group Accounts

Groups are collections of user accounts. Using groups is a convenient and efficient way to assign user rights and permissions to multiple users.

There are two fundamental types of groups in Windows 2000: security groups and distribution groups. *Security groups* are primarily used to assign permissions and user rights to multiple users. In addition, security groups can be used by some e-mail programs to send messages to the list of users who are members of the group.

Distribution groups are primarily used to send e-mail messages to a specified list of users. You can't assign permissions and user rights to distribution groups. Distribution groups are an important feature because some e-mail programs are unable to send e-mail to the list of users who are members of a security group. Lastly, distribution groups can't be created on the local computer — they can only be created in Active Directory.

Groups on the Local Computer

Groups on the local computer are primarily used to control access to resources on that computer. All groups on the local computer are security groups. There are two kinds of groups found on the local computer: local groups and built-in groups.

Local Groups

Local groups are groups that are created and maintained on an individual Windows 2000 computer (that is not a domain controller). Local groups can be created by members of the Administrators, Power Users, and Users groups.

Local groups are used to control access to resources on the local computer. In a typical configuration, a local group is assigned permissions to a specific resource, such as a shared folder or a shared printer. Then individual user accounts and groups are made members of this local group. The result

is that all members of the local group now have permissions to the shared resource on the local computer. Using local groups simplifies the administration of resources, because permissions can be assigned once to a local group, instead of separately to each user account.

Both local and domain user accounts can be members of a local group. In addition, built-in system groups on the local computer and global groups and universal groups from the domain can be members of a local group. Finally, a local group can't be a member of another group.

Built-in Groups

Built-in groups are groups with preset characteristics that are automatically created during the installation of Windows 2000. There are two kinds of built-in groups on a Windows 2000 computer that is not a domain controller: built-in local groups, and built-in special groups.

Built-in Local Groups

Built-in local groups are groups that have the rights and/or permissions that enable their members to perform specific tasks on the local computer. You can assign users to the built-in local groups that most closely match the tasks the users need to perform. If there isn't a built-in local group that has the rights or permissions needed to perform a specific task or access a specific resource, then you can create a local group and assign it the necessary rights or permissions to accomplish the task or access the resource.

You can assign rights and permissions to built-in local groups. In addition, you can make users members of (and remove users from) built-in local groups. (An exception is that you can't remove Administrator from the Administrators group.) Built-in local groups can be renamed, but they can't be deleted.

There are six built-in local groups that are automatically created during the installation of Windows 2000 on a nondomain controller:

- **Administrators:** Members of this group have full administrative rights and permissions to administer the local computer. This group initially contains the Administrator account, and, if the computer is a member of a domain, it contains that domain's Domain Admins global group.

- **Backup Operators:** Members of this group have permissions to back up and restore all files on the local computer, even if the user does not have permissions to all files. This group initially has no members.
- **Guests:** Members of this group can log on locally. Initially this group has no permissions to resources. This group initially contains the Guest account, which is disabled by default.
- **Power Users:** Members of this group can run applications, use local printers, and create local user and group accounts (and modify the users and groups they create). Members of this group can add users to and remove users from the Guests, Power Users, and Users groups. Members of this group can also share folders and printers. This group initially has no members.
- **Replicator:** This group, which supports directory replication processes, is included in Windows 2000 to provide backward compatibility with the Windows NT 4.0 Directory Replicator service. This group initially has no members.
- **Users:** Members of this group can run applications, create local groups (and manage the groups they create), and use local printers. This group initially contains the Authenticated Users and Interactive special groups, and, if the computer is a member of a domain, it contains that domain's Domain Users group. As new local user accounts are created, they are automatically made members of the built-in Users group.

Built-in Special Groups *Built-in special groups* are groups created by Windows 2000 that are used for specific purposes by the operating system. Special groups are sometimes called *system groups*.

You can assign user rights and permissions to special groups, and you can remove user rights and permissions from special groups. You can't assign users or groups to special groups. However, you can make a special group a member of a local group. You can't rename or delete special groups.

Membership in a special group is temporary, and is based solely on whether a specific set of membership requirements are met. A user is a member of a special group only for the time period in which the user meets the special group's membership requirements.

There are 12 built-in special groups on Windows 2000 computers that are not domain controllers:

- **Everyone:** Any user who accesses a Windows 2000 computer, either interactively or over-the-network, is considered a member of the Everyone special group. This includes all users accessing the computer using authorized user accounts, as well as users who are authenticated using an anonymous logon, such as a user who accesses a Web server over the network. If your network is connected to the Internet, over-the-network also means over the Internet. Because of this, Everyone means *everyone*. You should consider limiting the permissions assigned to the Everyone group.
- **Anonymous Logon:** Any user who accesses a Windows 2000 computer over-the-network (or over the Internet) by using an anonymous logon is considered a member of the Anonymous Logon special group.
- **Authenticated Users:** Any user who accesses a Windows 2000 computer, either interactively or over-the-network, by using an authorized user account is considered a member of the Authenticated Users special group.
- **Batch:** When a scheduled program or batch job logs on using a user account that has the "Log on as a batch job" user right, that user account is a member of the Batch special group.
- **Creator Owner:** A user who creates a file, folder, or print job is considered a member of the Creator Owner special group for that file, folder, or print job. The Creator Owner special group is used to assign permissions to creators of these objects. For example, by default the Creator Owner special group is assigned the Manage Documents permission to a printer when it is first created, so that creators of print jobs sent to this printer are able to manage their own print jobs.
- **Creator Group:** When a user of an Apple computer (or a user of a POSIX-compliant application) creates a file or folder, that user's primary group is considered a member of the Creator Group special group for that file or folder. The Creator Group special group is used to define the group ownership of the newly created file or folder.

- **Dialup:** Any user who accesses a Windows 2000 computer via a phone line, a Virtual Private Network (VPN), or a direct cable connection by using an authorized user account is considered a member of the Dialup special group.
- **Interactive:** Any user who physically sits at a computer and logs on locally to that Windows 2000 computer is a member of the Interactive special group. If you want to grant access to a resource on the local computer to users who log on locally to this computer, consider assigning the appropriate permissions to the Interactive group.
- **Network:** Any user who accesses resources on a Windows 2000 computer over-the-network is a member of the Network special group. If you want to grant access to a resource on the local computer to users who access this computer over-the-network, consider assigning the appropriate permissions to the Network group.
- **Service:** When a service logs on using a user account, that user account is a member of the Service special group.
- **System:** This special group is used by the Windows 2000 operating system. The System special group is not normally assigned any permissions to network resources.
- **Terminal Server User:** Any user who logs on to a Terminal Services session is a member of the Terminal Server User special group.

Groups in Active Directory

Groups in Active Directory are used to control access to network resources and to organize users who perform similar job tasks or have similar network access requirements.

There are three administrator-created kinds of groups in Active Directory: domain local groups, global groups, and universal groups. When you select one of these kinds of groups, the Windows 2000 user interface calls this selecting the *Group scope*. In addition to these three kinds of groups, there are built-in local, global, universal, and special groups in Active Directory.

Administrator-created groups in Active Directory can be either security groups or distribution groups. All of the built-in groups in Active Directory are security groups.

Domain Local Groups

Domain local groups are groups that are created and maintained in Active Directory on Windows 2000 domain controllers. Domain local groups are used to control access to resources located on any computer in a Windows 2000 domain.

In a typical configuration, a domain local group is assigned permissions to a specific resource, such as a shared folder or a shared printer. Then individual user accounts and groups are made members of this domain local group. The result is that all members of the domain local group now have permissions to the shared resource.

A domain local group can contain user accounts from its domain, and from other domains in the forest. A domain local group can contain other domain local groups from its own domain, and can also contain global and universal groups from any domain in the forest.

Global Groups

Global groups, like domain local groups, are created and maintained in Active Directory on Windows 2000 domain controllers. Global groups, however, are primarily used to organize users who perform similar tasks or have similar network access requirements.

In a typical configuration, user accounts of domain users who have similar job functions are placed in a global group. Then this global group is made a member of one or more domain local groups in any domain in the forest. Each of these domain local groups is assigned permissions to a specific shared resource. The result is that members of the global group now have permissions to the shared resource(s).

Here's an example of how global groups can be used in real life. Suppose that when the company's network was first installed, the administrator created user accounts, and placed these user accounts in various global groups depending on the users' job functions. Now, the network administrator wants to assign several users permissions to a shared printer on a Windows 2000 computer. The administrator creates a new domain local group and assigns this group permissions to the shared printer. Then the administrator selects the global groups that contain the user accounts that need access to this shared printer, and makes the global groups members of the new domain local group. The result is that all domain user accounts that are members of the selected global groups now have access to the shared printer. If access to all resources is managed in this way, when a new user is

created, the administrator need only make the user a member of the appropriate global group(s) in order for the user to have access to all network resources required to do his or her job.

The advantage of using global groups, then, is ease of administration — the network administrator can manage large numbers of users by placing them in a small number of global groups.

A global group can only contain user accounts and other global groups from its domain. Global groups can't contain domain local groups or universal groups from its domain, and can't contain user accounts or groups from any other domain.

Although it is not a preferred practice, you can assign user rights and permissions to global groups. Global groups can be assigned permissions to shared resources on any computer in the forest.

Universal Groups

Universal groups, like domain local groups and global groups, are created and maintained in Active Directory on Windows 2000 domain controllers. Universal groups, however, are used to organize users from *multiple domains* that perform similar job tasks or have similar network access requirements, and/or to control access to shared resources in *multiple domains*.

There's no one typical universal group configuration. For example, you can use a universal group as a "super" global group by placing users from multiple domains into the universal group, and then making the universal group a member of one or more domain local groups to which you have assigned permissions to shared resources. Or, you can use a universal group in much the same way as you'd use a domain local group, except that you can assign a universal group permission to a shared resource on any computer in the forest. In short, you can use universal groups just about any way you want to.

Universal groups provide significant advantages, but sometimes present significant challenges, too. The primary advantage of using universal groups is their open membership: user accounts, global groups, and universal groups from any domain in the forest can be members of a universal group. An addi- tional advantage of using universal groups is that universal groups can be assigned permissions to shared resources on any computer in the forest.

The main disadvantage of using universal groups is that they can cause potential network traffic problems. Here's how this can happen. When you first create a universal group, all of the group's members are listed in the global catalog. Then, each time you change the membership of a universal

group, the global catalog is updated, and this change is replicated to all global catalog servers on your network. If you have a large number of universal groups and change them frequently, this can cause significant amounts of replication traffic on your network.

Another challenge presented by universal groups is that they are *not* available if your Windows 2000 domain is operating in mixed-mode, that is, when you have both Windows 2000 domain controllers and Windows NT 4.1 backup domain controllers in your domain. Universal groups can only be used when your Windows 2000 domain is operating in native-mode.

Because of these challenges, you should only use universal groups when you need to organize users from multiple domains that perform similar job tasks or have similar network access requirements, or when you need to use a single group to control access to shared resources in multiple domains.

Built-in Groups on Domain Controllers

Built-in groups (in Active Directory) are security groups with preset characteristics that are automatically created during the installation of Active Directory. There are four kinds of built-in groups on Windows 2000 domain controllers: built-in local groups, built-in global groups, built-in universal groups, and built-in special groups.

Built-in Local Groups Built-in local groups on domain controllers are groups that are automatically created during the installation of Active Directory and stored in the `Builtin` folder. Built-in local groups have rights and/or permissions that enable their members to perform specific tasks in Active Directory and/or on Windows 2000 domain controllers in the domain.

You can assign rights and permissions to built-in local groups on domain controllers only for resources located in Active Directory and/or on domain controllers in the domain. You can also add members to and remove members from built-in local groups on domain controllers (except that you can't remove the Administrator account from the Administrators group).

Built-in local groups on domain controllers can contain user accounts from the domain and from other domains in the forest. In addition, built-in local groups on domain controllers can contain domain local groups from the domain, and global and universal groups from any domain in the forest.

Built-in local groups on domain controllers can't contain other built-it local groups. And, built-in local groups on domain controllers can't be members of any other groups.

There are nine built-in local groups that are automatically created during the installation of Active Directory on a Windows 2000 domain controller:

- **Account Operators:** Members of this group can create, delete, and modify domain user and group accounts in the domain, except that Account Operators can't modify the Administrator account and can't modify or change the membership of the Administrators, Account Operators, Backup Operators, Print Operators, or Server Operators groups. This group initially has no members.
- **Administrators:** Members of this group have full administrative rights and permissions to administer Active Directory (including all of its domain users, groups, and other objects) and all domain controllers in the domain. This group initially contains the Administrator account, the Domain Admins group, and the Enterprise Admins group.
- **Backup Operators:** Members of this group have permissions to back up and restore all files on all domain controllers in the domain, even if the user does not have permissions to all files. This group initially has no members.
- **Guests:** Members of this group have no initial rights or permissions. This group initially contains the Domain Guests group and the Guest account.
- **Pre-Windows 2000 Compatible Access:** Members of this group have the Read permission for all domain users and groups in the domain. This group initially has no members. The purpose of this group is to enable users of Windows NT 4.0 computers to log on to the domain. If you have Windows NT 4.0 computers in the domain, you should make the Everyone group a member of this group.
- **Print Operators:** Members of this group can create and manage printers on any domain controller in the domain. This group initially has no members.
- **Replicator:** This group, which supports directory replication processes, is included in Windows 2000 to provide backward compatibility with the Windows NT 4.0 Directory Replicator service. This group initially has no members.

- **Server Operators:** Members of this group have permissions to back up and restore files and folders on all domain controllers in the domain, and can share folders on any domain controller in the domain. This group initially has no members.
- **Users:** Members of this group have no initial rights or permissions. You can assign to this group rights and permissions that you want all domain users to have. This group initially contains the Authenticated Users, Domain Users, and Interactive groups. As new domain user accounts are created, they are automatically made members of the Domain Users group, which is a member of the built-in Users group.

Built-in Global and Universal Groups Built-in global and universal groups on domain controllers are automatically created during the installation of Active Directory and stored in the `Users` folder. Built-in global and universal groups are primarily used to group users by the types of administrative tasks they can perform in Active Directory and on all computers in the Windows 2000 domain.

Built-in global and universal groups on domain controllers have the same characteristics as administrator-created global and universal groups (which were covered earlier in this chapter).

There are numerous built-in global and universal groups. Below I've listed and described the most common ones:

- **Domain Admins:** Members of this global group have no initial rights or permissions. This group initially derives all of its rights and permissions from its membership in other groups. By default, this group is a member of the domain's built-in local Administrators group and the local built-in Administrators group on all computers that are members of the domain. As a result of this group's membership in other groups, members of Domain Admins can administer Active Directory and all computers in the domain. This group initially contains the Administrator account.
- **Domain Users:** Members of this global group have no initial rights or permissions. This group initially derives all of its rights and permissions from its membership in other groups. By default, this group is a member of the domain's built-in local Users group and the local built-in Users group on all computers that are members of the domain. This group initially contains all domain user accounts

created when Active Directory is installed, including the Administrator and Guest accounts. As new domain user accounts are created, they are automatically made members of Domain Users.

- **Domain Guests:** Members of this global group have no initial rights or permissions. This group initially derives all of its rights and permissions from its membership in the domain's built-in local Guests group. This group initially contains the Guest account.
- **Enterprise Admins:** Members of this universal group have no initial rights or permissions. This group initially derives all of its rights and permissions from its membership in the domain's built-in local Administrators group in each domain in the forest. As a result of this membership, members of Enterprise Admins can administer Active Directory throughout the forest and all domain controllers in the forest. If you have multiple domains in your forest, Windows 2000 only creates the Enterprise Admins group in the first domain in the forest. This group initially contains the Administrator account.
- **Schema Admins:** Members of this universal group can modify the Active Directory schema. If you have multiple domains in your forest, Windows 2000 only creates the Schema Admins group in the first domain in the forest. This group initially contains the Administrator account.

Built-in Special Groups Built-in special groups on domain controllers are automatically created during the installation of Active Directory. These built-in special groups are used for specific purposes by the operating system, and are sometimes called system groups.

All of the built-in special groups that exist on nondomain controllers are also present on Windows 2000 domain controllers. The built-in special groups on domain controllers have the same characteristics as the built-in special groups on nondomain controllers. (Built-in special groups on nondomain controllers are covered earlier in this chapter).