NEAR EAST UNIVERSITY

Faculty of Engineering

Department of Computer Engineering

Structure,Performance and Security Measures of Windows NT Based System

Graduation Project
COM - 400

Student: Ayhan TUNÇ

Supervisor: Dr.Halil ADAHAN

Nicosia 2004

Acknowledgements

First I want to thank Mr. Halil Adahan to be my advisor. Under his guidence, I succesfully overcome many difficulties and learn computer networks on Microsoft Windows 2000 operating systems performance measures, In each discussion, he explained my questions patiently, and I felt my quick progress from his advices. He always helps me a lot either in my study or my life. I asked him many questions in my subject and he always answered my questions quickly and in detail. His ideas has always been very usefull and he provided me with up to date information technology.

Special thanks to Ümit İlhan and Kaan Uyar for their practical advices. And thanks to Faculty of Engineering for having such a good computational and computer environment.

Finally, I want to thank my family, especially my parents. Without their endless support and love for me, I would never achieve my current position. I wish my mother father, my sisters and my brother live happily always.

Abstract

Modern today's applications involve the extensive use of Computer Networks and their applications. Computer Networks provide the use of many programs to be integrated to a singe environment where many users can run their applications without having to install the application on their machine. The simplifies the use of programs on end workstations.

Applications involve the use of Word Processing, multimedia such as Graphics programs, drawing tools as Autocad, and Internet based platforms such as Java, With integration of all applications integrated into the File Server the use and extensive multimedia programs will facilitate correct use and enforce security restrictions on program usage.

All facilities provided by computer networks make all Hardware as well as Software resources to be used worldwide.

Table of Contents

List of Figures

Chapter I: Networking Concepts

1.1 Microsoft 2000 Family

No matter where you are working, your computer will be easier to use and to manage, because Microsoft Windows 2000 Professional is more compatible and more powerful than any workstation you've used before. With Windows 2000 Professional, you have faster access to information, and you are able to accomplish tasks more quickly and easily.

Windows 2000 Advanced Server includes all the new features of Windows 2000 Server, and in addition offers enhanced memory support, support for additional processors, and clustering. Enhanced memory and processor support means your server applications can faster, providing better response for users on the network. Windows 2000 Professional makes it easier to:

1. Work with files.
2. Find information.
3. Personalize your computing environment.
4. Work on the Web.
5. Work remotely.

1.1.1 Manageability

You and your network administrators can work more efficiently now, because many of the most common computer-management tasks are automated and streamlined with Windows 2000 Professional. With Windows 2000, your workstation will be easier to:

1. Set up.
2. Administer.
3. Support.

1.1.1.1 Connection Manager and Administration Kit

Connection Manager is a versatile client dialer and connection software that you can customize by using the Connection Manager Administration Kit (CMAK) wizard.

1.1.1.2 Managing Event Logs

The Windows Server family supports two scripting environments: The command processor runs files that contain batch language commands. Batch language has limited capabilities, but earlier operating systems supported it and you may need to continue using existing batch files.

Windows Script Host (WSH) runs files that contain Microsoft Visual Basic Scripting Edition (VBScript) or Jscript commands. VBScript and Jscript provide all the capabilities of batch language plus many more. For example, scripts written in VBScript or Jscript can interact with Active Directory Service Interfaces (ADSI) to manage objects stored in Active Directory, and they can interact with Windows Management Instrumentation (WMI) to access system resources.

In addition to the two scripting languages, WSH provides two runtime programs: WScript.exe and CScript.exe. After you create a WSH script containing VBScript or Jscript commands, you use WScript or CScript to run the script. WScript runs the script as a Windows-based process and CScript runs the script as a console-based process.

Administrators often want to manage objects in Active Directory, such as organizational units, groups, and users. The following example script shows how you can use WSH, VBScript, and ADSI to create an Active Directory organizational unit, group, and user. After creating the three Active Directory objects in the current domain, the script assigns a password to the new user account, enables the user account, and adds the user account to the group.

1.1.1.3 Managing Applications on A Local Computer

Managing applications involves adding or removing applications, changing the configuration of an application, or scheduling an application to run at a specific time. You can use Add or Remove Programs in Control Panel to add or remove an application, modify an existing application, repair a damaged application, or add or remove components of Windows. You can use Task Scheduler or the command to schedule an application to run at a specific time.

You can also update your operating system and other Microsoft software by using Windows Update in Help and Support Center to download items such as security fixes, critical updates, the latest Help files, drivers, and Internet products.

Some of the most common tasks are adding, removing, or changing an application and scheduling an application to run at a specific time. You can also run an application at a scheduled time from the command line by using the command.

1.1.1.4 Managing Applications Remotely

You can manage the applications for an entire system from a remote location by using the remote administration features of the Windows Server family. With the Remote Desktops snap-in, you can access remote computers for installations, deletions, or modifications of applications on those computers. In addition, with Group Policy Software Installation, you can assign or publish applications. You assign an application to users or computers when you want everyone to have the application on his or her computer.

When the users log on to their computers, the application installs. You publish an application to users when you want the application to be available to the users who are managed by your Group Policy object. These users then determine when to install the application.

Some of the most common tasks are creating a new connection to a remote computer, assigning an application to users or computers, and publishing an application to users. For more information about other tasks for managing applications remotely.

1.1.1.5 Managing Directory Replication

Except for very small networks, directory data must reside in more than one place on the network to be equally useful to all users. Through replication, the Active Directory directory service maintains replicas of directory data on multiple domain controllers, ensuring directory availability and performance for all users. Active Directory relies on configuration information that you provide about sites, subnets, and site links to manage and optimize the process of replication.

Some of the most common tasks are creating sites, creating site links, and creating a subnet and associating it with a site.

1.1.1.6 Managing Disks and Volumes

Managing disks and volumes includes creating and formatting partitions, logical drives, and volumes; setting disk quotas to limit disk usage; defragmenting volumes to improve file-system performance; and checking for file-system errors and bad sectors on a hard disk. The Windows Server family provides many tools you can use to effectively manage disks and volumes on new or existing systems. These tools include Disk Management, Disk Defragmenter, disk quotas, and error checking.

Some of the most common tasks are creating partitions or logical drives, formatting basic volumes, extending basic volumes, and defragmenting volumes. You can also manage disks and volumes from the command line.

1.1.1.7 Managing Files and Folders

Managing files and folders includes storing and securing resources, making those resources available to network users, and managing changes in those resources. The Windows Server family provides many tools you can use to manage files and folders. These tools include Shared Folders, shadow copies of shared folders, Distributed File System (DFS), Encrypting File System (EFS), and Offline Files. When a folder is shared, users can connect to the folder over the network and gain access to the contents of the shared folder. With shadow copies of shared folders, users can view the contents of network folders as they existed at points of time in the past.

Some of the most common tasks are sharing a folder or drive, enabling shadow copies of shared folders, and changing settings for shadow copies of shared folders. You can also manage files and folders from the command line.

1.1.1.8 Server Management

Windows Server operating systems provide tools that you can use to manage servers from a remote location. These tools expand your flexibility because you can work as though you are physically present at each computer in your organization. The tools include the Remote Desktops snap-in, Active Directory Users and Computers, Web.

Some of the most common tasks are sharing a folder or drive, enabling shadow copies of shared folders, and changing settings for shadow copies of shared folders. You can also manage files and folders from the command line.

1.1.1.9 Managing Servers Remotely

Windows Server operating systems provide tools that you can use to manage servers from a remote location. These tools expand your flexibility because you can work as though you are physically present at each computer in your organization. The tools include the Remote Desktops snap-in, Active Directory Users and Computers, Web Interface for Remote Administration, and the Windows Server Administration Tools Pack. By understanding the advantages and security requirements of each tool, you can choose the most appropriate one for your remote administration and management tasks. For example, if you want to remotely manage a group of servers instead of just one server, Active Directory is the preferred tool to use. Each server that you want to manage remotely must be enabled for remote administration.

Some of the most common tasks are managing servers remotely by using the Remote Desktops snap-in and managing servers remotely by using Active Directory Users and Computers.

Administration Tools Pack to manage servers remotely from Windows XP Professional, see Windows Server Administration Tools Pack. For more information about using Web Interface for Remote Administration, see Using Web Interface for Remote Administration.

1.1.1.10 Managing Services

A service is an application type that runs in the background and is similar to a UNIX daemon application. Services typically provide access to key features such as file servers, Web servers, database servers, and other server-based applications to users, both locally and across the network.

Some of the most common tasks associated with managing services are starting and stopping a service, disabling a service for a hardware profile, and changing the startup method for a service.

To improve performance and security in the Windows Server family, several services have been disabled by default that were previously enabled. You should not change the startup method of a service unless you are sure you are choosing the appropriate startup method.

1.1.1.10.1 Windows Management Instrumentation

Windows Management Instrumentation command-line (WMIC) provides you with a simple command-line interface to Windows Management Instrumentation (WMI), so you can take advantage of WMI to manage computers running the Windows Server family of operating systems. WMIC interoperates with existing shells and utility commands, and can be easily extended by scripts or other administration-oriented applications.

1.1.1.11 Backing Up and Restoring Data

The Backup utility helps you protect data from accidental loss if your system's hardware or storage media fails. For example, you can use Backup to create a duplicate copy of the data on your hard disk and then archive the data on another storage device. The backup storage medium can be a logical drive such as your hard disk, a separate storage device such as a removable disk, or an entire library of disks or tapes organized and controlled by a robotic changer. If the original data on your hard disk is accidentally erased or overwritten, or becomes inaccessible because of a hard disk malfunction, you can easily restore the data from the archived copy.

You can use Backup to back up and restore data on FAT16, FAT32, or NTFS volumes. However, if you have backed up data from an NTFS volume, it is recommended that you restore the data to an NTFS volume of the same version to prevent losing data. Some file systems might not support all of the features of other file systems

1.1.1.12 Changing Group Membership

A group is a collection of users that you can use to simplify the administration of user permissions and rights. In addition, you can use a group to delegate administrative tasks, filter Group Policy settings, and create e-mail distribution lists. Users belonging to a particular group receive all the permissions and rights assigned to that group. By changing group memberships for a user, you can quickly change the resources to which that user has access, as well as the tasks delegated to the user and the Group Policy settings that apply to the user. You can change the membership of Active Directory groups to change users' permissions and rights within a domain or forest. You can also change the membership of local groups to change users' permissions and rights.

Some of the most common tasks are adding or removing members from Active Directory groups and adding or removing members from groups on a local computer. You can also use the command line to change group memberships, either in a domain or on a local computer.

1.1.1.13 Checking Event Logos

Depending on its server role, a computer running a Windows Server operating system records events in the following types of logs: application, security, directory service, File Replication service, and DNS server. You can use Event Viewer to monitor these logs and gather information about the hardware, software, and system problems on a computer. Event Viewer displays five types of events within each log: error, warning, information, success audit, and failure audit.

Some of the most common tasks are viewing an event log and connecting to another computer. You can also check event logs from the command line.

1.1.1.14 Creating Logon Scripts

You can use logon scripts to assign tasks that will be performed when a user logs on to a particular computer. The scripts can carry out operating system commands, set system environment variables, and call other scripts or executable programs. The Windows Server family supports two scripting environments: the command processor runs files containing batch language commands, and Windows Script Host (WSH) runs files containing Microsoft Visual Basic Scripting Edition (VBScript) or Jscript commands. You can use a text editor to create logon scripts. Some tasks commonly performed by logon scripts include:

- Mapping network drives.
- Installing and setting a user's default printer.
- Collecting computer system information.
- Updating virus signatures.
- Updating software.

The following example logon script contains VBScript commands that use Active Directory Service Interfaces (ADSI) to perform three common tasks based on a user's group membership: It maps the H: drive to the home directory of the user by calling the WSH Network object's MapNetworkDrive method in combination with the WSH Network object's UserName property.

It uses the ADSI IADsADSystemInfo object to obtain the current user's distinguished name, which in turn is used to connect to the corresponding user object in Active Directory. Once the connection is established, the list of groups the user is a member of is retrieved by using the user's memberOf attribute. The multivalued list of group names is joined into a single string by using VBScript's Join function to make it easier to search for target group names.

If the current user is a member of one of the three groups defined at the top of the script, then the script maps the user's G: drive to the group shared drive, and sets the user's default printer to be the group printer.

1.1.2 Compatibility

Windows 2000 Professional offers increased compatibility with different types of networks and with a wide array of legacy hardware and software. Windows 2000 also provides:

1. Improved driver support.
2. Increased support for new-generation hardware and multimedia technologies.
3. Integration of the new Euro currency symbol.

For all your computing needs, Windows 2000 Professional provides:

1. Industrial-strength reliability.
2. The highest level of security.
3. Powerful performance.

Windows 2000 Advanced Server includes all the new features of Windows 2000 Server, and in addition offers enhanced memory support, support for additional processors, and clustering. Enhanced memory and processor support means your server applications can faster, providing better response for users on the network.

1.2 File and Folder Structure

Almost all Windows 2000 tasks involve working with files and folders. The work you do with files and folders falls into three categories:

1. You can perform basic file and folder tasks, such as creating, deleting, copying, and moving files and folders, and more advanced tasks, such as changing file and folder properties and managing shared folders.
2. You can narrow the focus of your file and folder searches by including additional search criteria, such as the date, type, file size, or case sensitivity. You can also broaden the scope of your file searches by using wildcard characters, and specifying literal text or regular expressions.

3. You can secure files and folders using Windows 2000 Professional security features, such as user and group accounts, Group Policy, shared folder and printer permissions, auditing, and user rights. If you have an NTFS drive installed, you can set file and folder permissions and encrypt files and folders.

To open My Computer, double-click its icon on the desktop. To open a file or folder by using Windows Explorer, click **Start**, point to **Programs**, point to **Accessories**, click **Windows Explorer**, and then double-click the file or folder you want to open.

If the file you want to open is not associated with a particular program, you can select the program used to open the file by right-clicking the file, clicking **Open With**, and then selecting the name of the program.

You can use commands on the **View** menu to change the way files are displayed. You can also use the **View** tab in the **Folder Options** dialog box to change file and folder settings.

1.2.1 Shared Folder Characteristics

1. On the desktop, double-click **My Network Places**.
2. Locate and double-click the computer in which the shared folder is located.
3. Double-click the shared folder you want to open.

1.2.2 Network Connection

Network Connections provides connectivity between your computer and the Internet, a network, or another computer. With Network Connections, you can configure settings to reach local or remote network resources or functions.

Network Connections combines Microsoft Windows NT version 4.0 Dial-Up Networking with features that were formerly located in the Network Control Panel, such as network protocol and service configuration. Each connection in the Network Connections folder contains a set of features that creates a link between your computer and another computer or network.

By using Network Connections, performing a task, such as modifying a network protocol, is as easy as right-clicking a connection and then clicking Properties.

1.3 Networking Fundamentals

Networking lets you connect your computer to other computers or a private network. When you connect your computer to a network or another computer.

1. Gain access to files and folders on other computers.
2. Let other people gain access to your files and folders.
3. Use printers and other devices that are connected to other computers.
4. Let other people gain access to any printers or devices that are connected to your computer.

There are many different ways to connect your computer to another computer or a network. Using Windows 2000, you can connect your computer to:

1. Another computer using a direct cable connection.
2. A private network using a modem or an Integrated Services Digital Network (ISDN) adapter or a network adapter card.
3. A network using a virtual private network (VPN) connection.
4. Another computer by having another computer calls your computer.

You can make these connections and configure networking protocols and settings using Network and Dial-up Connections, which can be found in the Control Panel. You can also connect to bulletin board services, networks, and other computers using the Telnet or HyperTerminal utilities.

1.3.1 Virtual Private Network Connection Establishment

1. Open Network and Dial-up Connections.
2. Double-click **Make New Connection**, and then click **next**.
3. Click **Connect to a private network through the Internet**, and click **next**.
4. If you have already established a dial-up connection, do one of the following:

- If you need to establish a connection with your ISP or some other network before tunneling to your destination computer or network, click **automatically dials this initial connection**, click a connection in the list, and then click **next**.
- If you do not want to automatically dial an initial connection, click **do not dial the initial connection**, and then click **next**.

5. Type the host name or IP address of the computer or network to which you are connecting, and then click **Next**.
6. Do one of the following:
   - If you want this connection to be made available to all users on your network, click **for all users**, and then click **next**.
   - If you want to reserve the connection for your own use, click **only for myself**, and then click **next**.
7. If you want to let other computers access resources through this dial-up connection, select the **Enable Internet connection sharing for this connection** check box, and then click **next**.
8. Type a name for the connection, and then click **Finish**.

To open Network and Dial-up Connections, click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**. To make the connection available to all users, you must be logged on as Administrator or as a member of the Administrators group.

You can create multiple VPN connections by copying them in the Network and Dial-up Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple hosts, security options, and so on.

1.3.2 Direct Network Connection Establishment

1. Open Network and Dial-up Connections.
2. Double-click **Make New Connection**, and then click **next**.

3. Click **Connect directly to another computer**, click **next**, and then follow the instructions in the Network Connection wizard.

To open Network and Dial-up Connections, click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**. To create a direct network connection that acts as a host, you must be logged on as Administrator or be a member of the Administrators group. Guest direct network connections do not require administrator-level rights.

If you specify your connection as a host when you create it, the connection appears as **Incoming Connections** in the Network and Dial-up Connections folder. You can create multiple direct connections by copying them in the Network and Dial-up Connections folder. You can then rename the connections and modify connection settings. By doing so, you can easily create different connections to accommodate multiple ports, host computers, and so on.

Direct connections can bypass authentication requirements. This is useful for devices such as palmtop computers. You must configure this setting in the host incoming connection. For more information, see Related Topics. If you create a direct connection by using a serial (RS-232C) cable, the port that you select in the Network Connection wizard is enabled for connections that use a null modem.

If you are logged on to your computer as Administrator or a member of the Administrators group when you create a direct connection, you are presented with a list of connection devices to choose from that includes all of the parallel ports for the computer, infrared ports that are installed and enabled, and COM ports. If you are logged on as a user who is not a member of the Administrators group, and create a direct connection, the list of devices includes the parallel ports for the computer, infrared ports that are installed and enabled, and only the COM ports that are configured with null modems. If you need to use a COM port for a direct connection, ask your system administrator to configure one of the COM ports on your computer with a null modem by using Phone and Modem Options in Control Panel.

Users do not need to use direct connections to allow access to shared resources, such as files and printers, over a local area network. In order to enable shared access to resources on the local computer, you must enable file and print sharing, share the resources, and then set up the appropriate permissions.

1.3.3 Incoming Network Connection

1. Open Network and Dial-up Connections.
2. Double-click **Make New Connection**, and then click **next**.
3. Click **Accept incoming connections**, click **next**, and then follow the instructions in the Network Connection wizard.

To open Network and Dial-up Connections, click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**. If you make another incoming network connection, and you use the Network Connection wizard again, the existing incoming network connection is reconfigured. To create an incoming network connection, you must be a member of the Administrators group.

For large numbers of incoming connections on a computer running Windows 2000 Server that operates as part of a distributed network or as a domain controller, use Windows 2000 Server Routing and Remote Access to create a remote access server.

If your incoming connection and Fax Service have problems working together (for example, you cannot receive incoming connection calls on a device enabled for Fax receive), the modem may not support adaptive answer. Check your modem documentation to verify that you need to disable Fax receives for that device to accept incoming connections.

If you connect to a computer running Windows 2000 Professional or stand-alone Windows 2000 Server that is configured for incoming connections, and you are running Windows 95 or Windows 98 and want to log on to the computer by using a local user account, you can use your Windows 95 or Windows 98 user name, domain, and password.

When you connect, the computer running Windows 2000 replaces the Windows 95 or Windows 98 domain name with the local computer name when you provide your user name and password authentication information. Incoming connections are only used for dial-up, VPN, or direct connection clients.

1.3.3.1 Network Component Addition

1. Open Network and Dial-up Connections.
2. Right-click the connection to which you want to add a network component, and then click **Properties**.
3. Do one of the following:
    - If this is a local area connection, click **Install**.
    - If this is a dial-up, VPN, or incoming connection, on the **Networking** tab, click **Install**.
4. In the **Select Network Component Type** dialog box, click Client, **Service**, or **Protocol**, and then click **Add**.
5. Do one of the following:
    - If you do not have an installation disk for the component, click the appropriate client, service, or protocol, and then click **OK**.
    - If you have an installation disk for the component, click the appropriate client, service, or protocol, click **Have Disk**, insert the installation disk into the selected drive, and then click **OK**.

To open Network and Dial-up Connections, click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**. You should only install the network components that you need, for the following reasons:

1. Network performance is enhanced and network traffic is reduced when only the required protocols and clients are installed.

2. If Windows 2000 encounters a problem with a network or dial-up connection, it attempts to establish connectivity by using every network protocol that is installed and enabled. By only installing and enabling the protocols that your system can use, Windows 2000 does not attempt to connect with protocols it cannot use, and returns status information to you more efficiently.
3. Excessive services can hinder performance on your local computer.

1.3.3.2 Monitoring Network Traffic

As an administrator, you need to monitor and detect problems with traffic on your network. With Network Monitor, you can gather information about the network traffic that flows to and from the network adapter of the computer on which it is installed. Once you capture the information, you can use Network Monitor to analyze the information, diagnose problem traffic patterns, and devise strategies to prevent future network traffic problems.

Some of the most common tasks are installing Network Monitor, specifying data frame patterns to capture, capturing network frames, and viewing a specific frame. You can also monitor network traffic from the command line.

1.3.3.3 Monitoring Security – Related Events

It is important to monitor security related events so that you can find out if changes are made to security policies or other objects, who made the changes, and when they were made. You can monitor security-related events by setting up an auditing policy. An auditing policy is made up of all of the auditing settings that you configure for individual security event categories.

Some of the most common tasks for monitoring security-related events are defining or modifying auditing policy settings for an event category on your local computer, defining or modifying auditing policy settings for an event category across your organization, and viewing the security log.

1.3.3.4 Monitoring Server Performance

In today's business environment, administrators need to ensure their computer systems are efficient and reliable. To optimize the performance of your servers, you need the data supplied by performance monitoring. This topic describes the most common tasks associated with monitoring the performance of a simple server configuration, one that has a few client workstations connected to a single facility that contains one or more servers. To monitor the performance of a simple server configuration, you need to collect three different types of performance data over a period of time:

General performance data: Information that can help you identify short-term trends such as memory leaks. After a month or two of data collection, you can average the results and save them in a more compact format. This archived data can assist you in capacity planning as your business grows, and later help you to evaluate the effectiveness of your plan.

Baseline performance data: Information that can help you discover changes that occur slowly, over time. By comparing the current state of your system with historical data, you can troubleshoot and tune your system. Because this information is collected only periodically, there is no need to compress it for storage.

Data for service level reports: Information that can help you ensure that your system meets a certain service or performance level, and which you will likely present to decision makers who are not performance analysts. How often you collect and maintain this data depends on your specific business needs.

To collect all three types of data, you can use Performance Logs and Alerts to create a counter log. You can also collect this information from the command line. You can then run the log over time, either manually or with automated scheduling. You can customize a counter log by adding objects and adding counters.

1.3.3.5 TCP/IP Installation

1. Open Network and Dial-up Connections.
2. Right-click the network connection for which you want to install and enable TCP/IP, and then click **Properties**.
3. On the **General** tab (for a local area connection) or the **Networking** tab (all other connections), if **Internet Protocol (TCP/IP)** is not in the list of installed components, then do the following:
    a. Click **Install**.
    b. Click **Protocol**, and then click **Add**.
    c. In the **Select Network Protocol** dialog box, click **Internet Protocol (TCP/IP)**, and then click **OK**.
4. Verify that the **Internet Protocol (TCP/IP)** check box is selected, and then click **OK**.

TCP/IP is installed as the default network protocol if network adapter hardware was detected during Windows 2000 Setup. You only need to follow these instructions if the TCP/IP default selection was overridden during Setup.

1.3.3.6 Setting up TCP/IP

By default, computers running Windows Server operating systems are configured as Dynamic Host Configuration Protocol (DHCP) clients, which means they can accept leases from any available DHCP server. TCP/IP computers (hosts) on a DHCP network automatically obtain TCP/IP configurations from DHCP servers that are located elsewhere on that network. This is the most common scenario for mid-sized and large TCP/IP networks.

In certain cases, you will need to configure TCP/IP for static addressing. Configuring TCP/IP for static addressing is the process of manually assigning unique, permanent addresses to network devices, and providing static address information about other network resources (for example, DHCP servers, routers, and default gateways). If you are setting up a computer to run on a network that does not have a DHCP server, but that does provide Windows Internet Name Service (WINS), you will need to configure TCP/IP to use WINS. For more information about WINS, see WINS.

Some of the most common tasks for setting up TCP/IP are configuring TCP/IP for static addressing, and configuring TCP/IP to use WINS. You can also set up TCP/IP from the command line. For more information about other tasks for setting up TCP/IP,

1.3.3.7 Telephony

This implementation of Telephony Application Programming Interface (TAPI) unifies Internet Protocol (IP) and traditional public switched telephone network (PSTN) telephony to support programs that effectively work the same way over intranets, the Internet, and traditional networks. This section provides information for configuring and managing both IP and PSTN telephony servers and clients.

1.3.3.8 DHCP

Dynamic Host Configuration Protocol (DHCP) is an IP standard designed to reduce the complexity of administering address configurations by using a server computer to centrally manage IP addresses and other related configuration details used on your network. The Microsoft Windows Server family provides the DHCP service, which enables the server computer to perform as a DHCP server and configure DHCP-enabled client computers on your network as described in the current DHCP draft standard, RFC 2131.

DHCP includes Multicast Address Dynamic Client Assignment Protocol (MADCAP) which is used to perform multicast address allocation. When registered clients are dynamically assigned IP addresses through MADCAP, they can participate efficiently in the data stream process, such as for real-time video or audio network transmissions.

1.3.3.9 Installing A DHCP Server

All computers on a TCP/IP network must have an IP address in order for the network to properly function. You can configure IP addresses manually at each computer, or you can install a DHCP server that automatically assigns IP address leases to each client computer on the network. Most client operating systems seek an IP address lease by default, so no configuration on the client computer is necessary to implement a DHCP-enabled network.

Before a DHCP server can start leasing IP addresses to client computers, you must create and activate a scope. A scope is a range of possible IP addresses for a network. Before you create a scope, make sure that the IP address range you want to use provides enough IP addresses for all of the computers on your network. Also, determine whether any devices on your network, such as DNS servers, WINS servers, or legacy printers, will need to use static IP addresses.

If there are devices that need static IP addresses, create an exclusion range of IP addresses at the beginning of the IP address range. An exclusion range is a group of IP addresses that the DHCP server will not lease to client computers. Once the exclusion range has been defined, you can assign all statically configured devices an IP address from the exclusion range. Some of the most common tasks are installing a DHCP server, creating a scope, and activating a scope. You can also perform these tasks from the command line.

1.3.3.10 Deploying and Upgrading Software

You can deploy and upgrade software to remote computers in managed environments by using Group Policy Software Installation to assign Windows Installer packages. Windows Installer packages are deployed and managed within a Group Policy object, which is in turn associated with a particular Active Directory container—either a site, a domain, or an organizational unit. You can use Add or Remove Programs in Control Panel to install, upgrade, or manage an application on a local computer. You can also use Remote Desktop Connection to install or upgrade an application by using Add or Remove Programs on a remote computer.

Some of the most common tasks are deploying software to remote computers, upgrading software on remote computers, and installing or upgrading software on a local computer.

1.3.3.11 WINS

Windows Internet Name Service (WINS) provides a dynamic replicated database service that can register and resolve NetBIOS names to IP addresses used on your network. The Microsoft Windows Server family provides WINS, which enables the server computer to act as a NetBIOS name server and register and resolve names for WINS-enabled client computers on your network as described in the NetBIOS over TCP/IP standards.

1.3.3.12 Routing and Remote Access

The Routing and Remote Access service in the Microsoft Windows Server family provides: Multiprotocol LAN-to-LAN, LAN-to-WAN, virtual private network (VPN), and network address translation (NAT) routing services.

1.3.3.13 Modems

Modem support is provided by the TAPI Unimodem 5 Service Provider. It enables you to install and use almost any modem available today. Modems are used by programs such as HyperTerminal, Network Connections, and Fax, as well as by many products, including software required to gain access to Internet service providers.

1.3.4 Wireless Networking

Wireless technology makes it possible for you to use a wide range of devices to access data from anywhere in the world. Wireless networks reduce or eliminate the high cost of laying expensive fiber and cabling and provide backup functionality for wired networks. To ensure that wireless networks and devices are compatible, cost-effective, and secure, organizations and special interest groups are working to develop standards for wireless communications.

The Microsoft Windows Server family provides support for infrared communication and for 802.11 wireless networks. In Windows Server , Standard Edition, support is provided for infrared communication. 802.11 wireless networking is supported in all versions of the Windows Server family.

1.3.5 Permanent Virtual Connection Using ATM

1. Open Network and Dial-up Connections.
2. Click the ATM connection that corresponds to the ATM network adapter installed on this computer for which you want to create a permanent virtual circuit (PVC).
3. Click **File**, and then click **Properties**.
4. In the list of network components used in this connection, select **ATM Call Manager**, and then click **Properties**.
5. In **ATM Call Manager** properties, click **Add**.

6. Review and modify PVC settings as needed:
    - For **Name**, you can either use the default unspecified PVC name or type a name. Both are used only for your reference.
    - For **Virtual path ID**, you can either use the default path of **0** or type a number that should be used to identify the virtual path for the connection.
    - For **Virtual circuit ID**, type a number that identifies the virtual circuit within the specified virtual path for the connection.
    - In **Application type**, select the type of application or use for this permanent virtual connection. If you configured your IP/ATM connection for PVCs only, you must select the application type **ATM ARP** for this PVC.
7. If needed, click **advanced** to configure any settings that provide call or answer matching criteria for the PVC or that specify a quality of service for use with the PVC.

To open Network and Dial-up Connections, click **Start**, point to **Settings**, and then click **Network and Dial-up Connections**.

1.3.6 SNMP Service Management

Open the Windows Components wizard.

1. In **Components**, click **Management and Monitoring Tools** (but do not select or clear its check box), and then click **Details**.
2. Select **Simple Network Management Protocol** check box, and click **OK**.
3. Click **Next**.

To open the Windows Components wizard, click **Start**, point to **Settings**, click **Control Panel**, double-click **Add/Remove Programs**, and then click **Add/Remove Windows Components**. Certain Windows components require configuration before they can be used. If you installed one or more of these components, but did not configure them, when you click **Add/Remove Windows Components**, a list of components that need to be configured is displayed. To start the Windows Components wizard, click **Components**.

You must be logged on as an administrator or a member of the Administrators group in order to complete this procedure. If your computer is connected to a network, network policy settings might also prevent you from completing this procedure. SNMP starts automatically after installation.

1.3.7 IPSec Policy

Open **Network and Dial-up Connections**.

1. Click **Local Area Connection**, and on the **File** menu, click **Properties**.
2. In the **Local Area Connection Properties** dialog box, under **Components checked are used by this connection**, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
3. Click **Advanced**, and then click the **Options** tab.
4. Under **Optional settings**, click **IP security**, and then click **Properties**.
5. Click **Use this IP security policy**, and then select the IPSec policy you want from the drop-down list.

You must be a member of the Administrators group to set Internet Protocol security (IPSec) policies. If the computer participates in a Windows 2000 domain, the computer may receive the IPSec policy from Active Directory, overriding the local IPSec policy. In this case, the options are disabled and you cannot change them from the local computer.

To open Network and Dial-up Connections, click **Start**, point to **Settings**, click **Control Panel**, and then double-click **Network and Dial-up Connections**. There are three predefined security policies: Client (Respond Only), Server (Request Security), and Secure Server (Require Security):

A1-) Internet Protocol security (IPSec) is a framework of open standards for ensuring private, secure communications over Internet Protocol (IP) networks, through the use of cryptographic security services.

Activating the Client (Respond Only) policy will not secure traffic unless the destination computer requests it. A server policy may need to be customized to work transparently with some programs and networks.

2.4 DNS

The Domain Name System (DNS) is an Internet and TCP/IP standard name service. The DNS service enables client computers on your network to register and resolve DNS domain names. These names are used to find and access resources offered by other computers on your network or other networks, such as the Internet.

DNS is an abbreviation for Domain Name System, a system for naming computers and network services that is organized into a hierarchy of domains. DNS naming is used in TCP/IP networks, such as the Internet, to locate computers and services through user-friendly names. When a user enters a DNS name in an application, DNS services can resolve the name to other information associated with the name, such as an IP address.

For example, most users prefer a friendly name such as example.microsoft.com to locate a computer such as a mail or Web server on a network. A friendly name can be easier to learn and remember. However, computers communicate over a network by using numeric addresses. To make use of network resources easier, name services such as DNS provide a way to map the user-friendly name for a computer or service to its numeric address. If you have ever used a Web browser, you have used DNS.

In this example, a client computer queries a server, asking for the IP address of a computer configured to use host-a.example.microsoft.com as its DNS domain name. Because the server is able to answer the query based on its local database, it replies with an answer containing the requested information, which is a host (A) resource record that contains the IP address information for host-a.example.microsoft.com.

The example shows a simple DNS query between a single client and server. In practice, DNS queries can be more involved than this and include additional steps not shown here.

Domain controllers provide network users and computers with the Active Directory service, which stores and replicates directory data and manages user interactions with the domain, including user logon processes, authentication, and directory searches. Every domain must contain at least one domain controller. You install a domain controller by installing Active Directory on any member or stand-alone server.

When you install the first domain controller in your organization, you are creating the first domain, also called the root domain and the first forest. You can add additional domain controllers to an existing domain to provide fault tolerance, improve service availability, and balance the load of existing domain controllers.

You can also install a domain controller to create a new child domain or new domain tree. Create a new domain when you want a new domain that shares a contiguous namespace with one or more domains. This means that the name of the new domain contains the full name of the parent domain.

For example, sales.microsoft.com would be a child domain of microsoft.com. Create a new domain tree only when you need a domain whose Domain Name System (DNS) namespace is not related to the other domains in the forest. This means that the name of the new domain tree's root domain (and all of its children) does not contain the full name of the parent domain. A forest can contain one or more domain trees. Before installing a new domain controller, you will need to consider pre-Windows 2000 compatible security levels and identify the DNS name of the domain.

The most commonly performed tasks when installing a domain controller are creating a new domain in a new forest, creating a new child domain in an existing domain tree, creating a new domain tree in an existing forest, and adding a domain controller to an existing domain.

1.4.1 DNS Domain Names

A DNS domain namespace, which specifies a structured hierarchy of domains used to organize names. Resource records, which map DNS domain names to a specific type of resource information for use when the name is registered or resolved in the namespace.

DNS servers, which store and answer, name queries for resource records. DNS clients, also known as resolves, which query servers to look up and resolve names to a type of resource record specified in the query.

1.4.2 The DNS Domain Namespace

The DNS domain namespace, as shown in the following figure, is based on the concept of a tree of named domains. Each level of the tree can represent either a branch or a leaf of the tree. A branch is a level where more than one name is used to identify a collection of named resources. A leaf represents a single name used once at that level to indicate a specific resource.

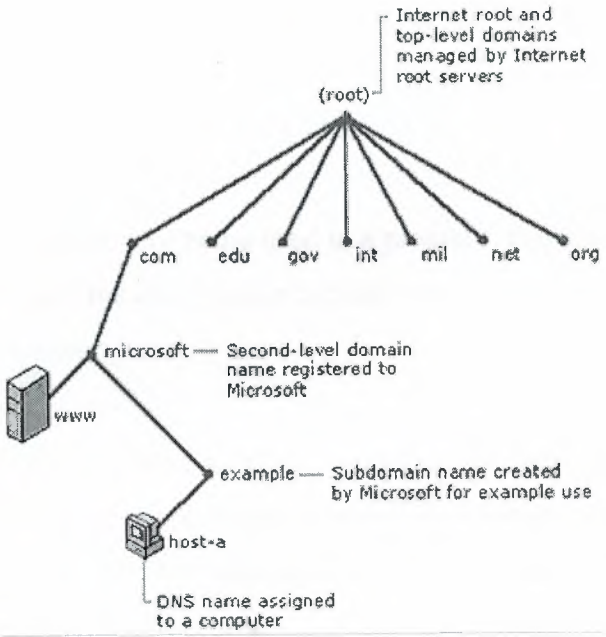

Figure 1.4.2: DNS Domain Namespace.

The previous graphic shows how Microsoft is assigned authority by the Internet root servers for its own part of the DNS domain namespace tree on the Internet. DNS clients and servers use queries as the fundamental method of resolving names in the tree to specific types of resource information. This information is provided by DNS servers in query responses to DNS clients, who then extract the information and pass it to a requesting program for resolving the queried name.

In the process of resolving a name, keep in mind that DNS servers often function as DNS clients, querying other servers in order to fully resolve a queried name.

Any DNS domain name used in the tree is technically a domain. Most DNS discussions, however, identify names in one of five ways, based on the level and the way a name is commonly used. For example, the DNS domain name registered to Microsoft (microsoft.com.) is known as a second-level domain. This is because the name has two parts (known as labels) that indicate it is located two levels below the root or top of the tree. Most DNS domain names have two or more labels, each of which indicates a new level in the tree. Periods are used in names to separate labels.

1.4.3 DNS Query

When a DNS client needs to look up a name used in a program, it queries DNS servers to resolve the name. Each query message the client sends contains three pieces of information, specifying a question for the server to answer:

A specified DNS domain name, stated as a fully qualified domain name (FQDN). A specified query type, which can either specify a resource record by type or a specialized type of query operation a specified class for the DNS domain name. For Windows DNS servers, this should always be specified as the Internet (IN) class.

DNS queries resolve in a number of different ways. A client can sometimes answer a query locally using cached information obtained from a previous query.

The DNS server can use its own cache of resource record information to answer a query. A DNS server can also query or contact other DNS servers on behalf of the requesting client to fully resolve the name, and then send an answer back to the client. This process is known as recursion. In addition, the client itself can attempt to contact additional DNS servers to resolve a name. When a client does so, it uses separate and additional queries based on referral answers from servers. This process is known as iteration. In general, the DNS query process occurs in two parts:

1. A name query begins at a client computer and is passed to a resolve, the DNS Client service, for resolution.
2. When the query cannot be resolved locally, DNS servers can be queried as needed to resolve the name.

1.4.3.1 The Local Resolver: Part I

The following graphic shows an overview of the complete DNS query process.

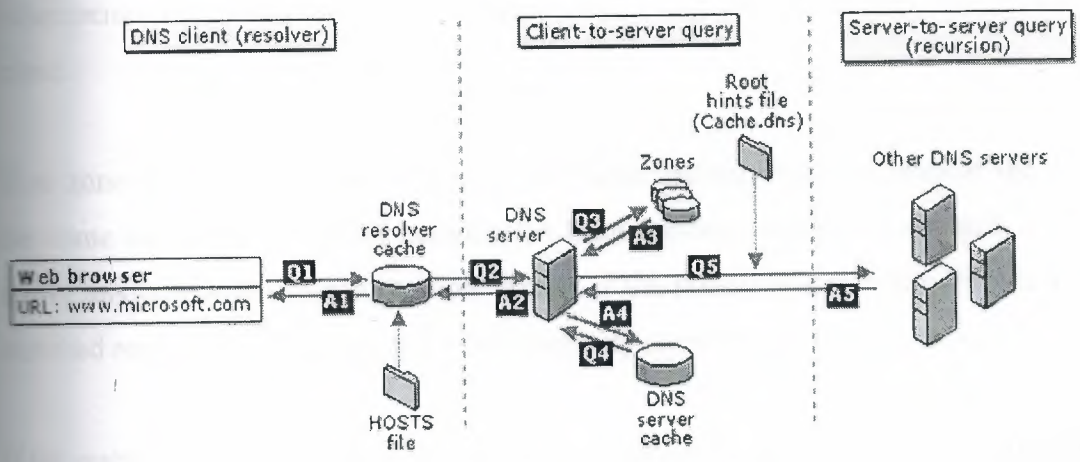

Figure 1.4.3.1 DNS Query Process.

As shown in the initial steps of the query process; a DNS domain name is used in a program on the local computer. The request is then passed to the DNS Client service for resolution using locally cached information. If the queried name can be resolved, the query is answered and the process is completed.

The local resolver cache can include name information obtained from two possible sources:

1. If a Hosts file is configured locally, any host name-to-address mappings from that file are preloaded into the cache when the DNS Client service is started.
2. Resource records obtained in answered responses from previous DNS queries are added to the cache and kept for a period of time.

If the query does not match an entry in the cache, the resolution process continues with the client querying a DNS server to resolve the name.

1.4.3.2 Querying A DNS Server: Part II

The actual server used during the initial client/server query part of the process is selected from a global list.

When the DNS server receives a query, it first checks to see if it can answer the query authoritatively based on resource record information contained in a locally configured zone on the server. If the queried name matches a corresponding resource record in local zone information, the server answers authoritatively, using this information to resolve the queried name.

If no zone information exists for the queried name, the server then checks to see if it can resolve the name using locally cached information from previous queries. If a match is found here, the server answers with this information. Again, if the preferred server can answer with a positive matched response from its cache to the requesting client, the query is completed.

If the queried name does not find a matched answer at its preferred server - either from its cache or zone information - the query process can continue, using recursion to fully resolve the name. This involves assistance from other DNS servers to help resolve the name. By default, the DNS Client service asks the server to use a process of recursion to fully resolve names on behalf of the the client before returning an answer.

In most cases, the DNS server is configured, by default, to support the recursion process. In order for the DNS server to do recursion properly, it first needs some helpful contact information about other DNS servers in the DNS domain namespace.

This information is provided in the form of *root hints*, a list of preliminary resource records that can be used by the DNS service to locate other DNS servers that are authoritative for the root of the DNS domain namespace tree. Root servers are authoritative for the domain root and top-level domains in the DNS domain namespace tree.

By using root hints to find root servers, a DNS server is able to complete the use of recursion. In theory, this process enables any DNS server to locate the servers that are authoritative for any other DNS domain name used at any level in the namespace tree.

For example, consider the use of the recursion process to locate the name "host-b.example.microsoft.com." when the client queries a single DNS server. The process occurs when a DNS server and client are first started and have no locally cached information available to help resolve a name query. It assumes that the name queried by the client is for a domain name of which the server has no local knowledge, based on its configured zones.

First, the preferred server parses the full name and determines that it needs the location of the server that is authoritative for the top-level domain, "com". It then uses an iterative query to the "com" DNS server to obtain a referral to the "microsoft.com" server. Next, a referral answer comes from the "microsoft.com" server to the DNS server for "example.microsoft.com".

Finally, the "example.microsoft.com." server is contacted. Because this server contains the queried name as part of its configured zones, it responds authoritatively back to the original server that initiated recursion. When the original server receives the response indicating that an authoritative answer was obtained to the requested query, it forwards this answer back to the requesting client and the recursive query process is completed. Although the recursive query process can be resource-intensive when performed as described above, it has some performance advantages for the DNS server.

For example, during the recursion process, the DNS server performing the recursive lookup obtains information about the DNS domain namespace. This information is cached by the server and can be used again to help speed the answering of subsequent queries that use or match it. Over time, this cached information can grow to occupy a significant portion of server memory resources, although it is cleared whenever the DNS service is cycled on and off.

1.4.4 Query Responses

The previous discussion of DNS queries assumes that the process ends with a positive response returned to the client. However, queries can return other answers as well. These are the most common:

1. An authoritative answer.
2. A positive answer.
3. A referral answer.
4. A negative answer.

An authoritative answer is a positive answer returned to the client and delivered with the authority bit set in the DNS message to indicate the answer was obtained from a server with direct authority for the queried name. A positive response can consist of the queried RR or a list of RRs (also known as an RRset) that fits the queried DNS domain name and record type specified in the query message.

A referral answer contains additional resource records not specified by name or type in the query. This type of answer is returned to the client if the recursion process is not supported. The records are meant to act as helpful reference answers that the client can use to continue the query using iteration.

A referral answer contains additional data such as resource records (RRs) that are other than the type queried. For example, if the queried host name was "www" and no A RRs for this name were found in this zone but a CNAME RR for "www" was found instead, the DNS server can include that information when responding to the client.

If the client is able to use iteration, it can make additional queries using the referral information in an attempt to fully resolve the name for it.

A negative response from the server can indicate that one of two possible results was encountered while the server attempted to process and recursively resolve the query fully and authoritatively:

1. An authoritative server reported that the queried name does not exist in the DNS namespace.
2. An authoritative server reported that the queried name exists but no records of the specified type exist for that name.

The resolver passes the results of the query, in the form of either a positive or negative response, back to the requesting program and caches the response.

1. If the resultant answer to a query is too long to be sent and resolved in a single UDP message packet, the DNS server can initiate a failover response over TCP port 53 to answer the client fully in a TCP connected session.
2. Disabling the use of recursion on a DNS server is generally done when DNS clients are being limited to resolving names to a specific DNS server, such as one located on your intranet. Recursion might also be disabled when the DNS server is incapable of resolving external DNS names, and clients are expected to fail over to another DNS server for resolution of these names.
3. For Windows 2000 Server, you can disable the use of recursion for DNS servers as needed by configuring in **advanced** properties in the DNS console on the applicable server.
4. By default, Windows 2000 DNS servers use several default timings when performing a recursive query and contacting other DNS servers. These are:
    - A recursion retry interval of 3 seconds. This is the length of time the DNS service waits before retrying a query made during a recursive lookup.

- A recursion time-out interval of 15 seconds. This is the length of time the DNS service waits before failing a recursive lookup that has been retried.

Under most circumstances, these parameters do not need adjustment. However, if you are using recursive lookups over a slow-speed WAN link, you might be able to improve server performance and query completion by making slight adjustments to the settings.

1.4.4.1 Caching

As DNS servers process client queries using recursion or iteration, they discover and acquire a significant store of information about the DNS namespace. This information is then cached by the server. Caching provides a way to speed the performance of DNS resolution for subsequent queries of popular names, while substantially reducing DNS-related query traffic on the network.

As DNS servers make recursive queries on behalf of clients, they temporarily cache resource records (RRs). Cached RRs contain information obtained from DNS servers that are authoritative for DNS domain names learned while making iterative queries to search and fully answer a recursive query performed on behalf of a client. Later, when other clients place new queries that request RR information matching cached RRs, the DNS server can use the cached RR information to answer them.

When information is cached, a Time-To-Live (TTL) value applies to all cached RRs. As long as the TTL for a cached RR does not expire, a DNS server can continue to cache and use the RR again when answering queries by its clients that match these RRs. Caching TTL values used by RRs in most zone configurations are assigned the **Minimum (default) TTL** which is set used in the zone's start of authority (SOA) resource record. By default, the minimum TTL is 3,600 seconds (1 hour) but can be adjusted or, if needed, individual caching TTLs can be set at each RR.

By default, Windows 2000 DNS servers use a root hints file, Cache.dns, that is stored in the %SystemRoot%\System32\Dns folder on the server computer. The contents of this file are preloaded into server memory when the service is started and contain pointer information to root servers for the DNS namespace where you are operating DNS servers.

1.4.4.2 Reverse Lookup

In most DNS lookups, clients typically perform a forward lookup, which is a search based on the DNS name of another computer as stored in an address (A) resource record. This type of query expects an IP address as the resource data for the answered response.

DNS also provides a reverse lookup process, enabling clients to use a known IP address during a name query and look up a computer name based on its address. A reverse lookup takes the form of a question, such as "Can you tell me the DNS name of the computer that uses the IP address 192.168.1.20?"

DNS was not originally designed to support this type of query. One problem for supporting the reverse query process is the difference in how the DNS namespace organizes and indexes names and how IP addresses are assigned. If the only method to answer the previous question was to search in all domains in the DNS namespace, a reverse query would take too long and require too much processing to be useful.

To solve this problem, a special domain, the in-addr.arpa domain, was defined in the DNS standards and reserved in the Internet DNS namespace to provide a practical and reliable way to perform reverse queries.

To create the reverse namespace, sub domains within the in-addr.arpa domain are formed using the reverse ordering of the numbers in the dotted-decimal notation of IP addresses.This reversed ordering of the domains for each octet value is needed because, unlike DNS names, when IP addresses are read from left to right, they are interpreted in the opposite manner.

When an IP address is read from left to right, it is viewed from its most generalized information (an IP network address) in the first part of the address to the more specific information (an IP host address) contained in the last octets.

For this reason, the order of IP address octets must be reversed when building the in-addr.arpa domain tree. With this arrangement, administering lower limbs of the DNS in-addr.arpa tree can be given to companies as they are assigned a specific or limited set of IP addresses within the Internet-defined address classes.

Finally, the in-addr.arpa domain tree, as built into DNS, requires that an additional resource record (RR) type - the pointer (PTR) RR - be defined. This RR is used to create a mapping in the reverse lookup zone that typically corresponds to a host (A) named RR for the DNS computer name of a host in its forward lookup zone.

The in-addr.arpa domain applies for use in all TCP/IP networks that are based on Internet Protocol version 4 (IPv4) addressing. The New Zone wizard automatically assumes that you are using this domain when creating a new reverse lookup zone.

If you are installing DNS and configuring reverse lookup zones for an Internet Protocol version 6 (IPv6) network, you can specify an exact name in the New Zone wizard. This will permit you to create reverse lookup zones in the DNS console that can be used to support IPv6 networks, which uses a different special domain name, the ip6.int domain.

1.5 Reverse Query

The following graphic shows an example of a reverse query initiated by a DNS client (host-b) to learn the name of another host (host-a) based on its IP address, 192.168.1.20.

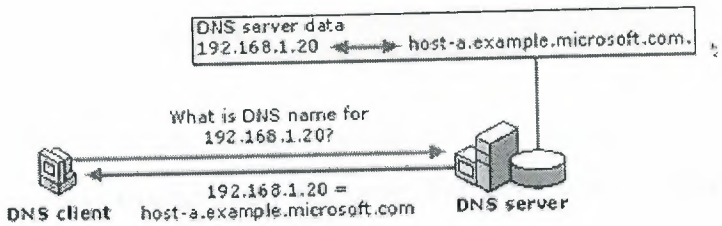

Figure 1.5: Reverse Query.

The reverse query process as shown in this graphic occurs in the following steps:

1. The client, "host-b", queries the DNS server for a pointer (PTR) RR that maps to the IP address of 192.168.1.20 for "host-a". Because the query is for PTR records, the resolver reverses the address and appends the in-addr.arpa domain to the end of the reverse address. This forms the fully qualified domain name ("20.1.168.192.in-addr.arpa.") for which to be searched in a reverse lookup zone.

2. Once located, the authoritative DNS server for "20.1.168.192.in-addr.arpa" can respond with the PTR record information. This includes the DNS domain name for "host-a", completing the reverse lookup process.

Keep in mind that if the queried reverse name is not answerable from the DNS server, normal DNS resolution (either recursion or iteration) can be used to locate a DNS server that is authoritative for the reverse lookup zone and that contains the queried name. In this sense, the name resolution process used in a reverse lookup is identical to that of a forward lookup.

For Windows 2000 Server, the DNS snap-in provides a means for you to configure a sub netted reverse lookup "classless" zone when the **advanced** view is selected. This allows you to configure a zone in the in-addr.arpa domain for a limited set of assigned IP addresses where a no default IP subnet mask is used with those addresses.

1.6 Inverse Queries

Inverse queries are an outdated practice, originally proposed as part of the DNS standard to look up a host name based on its IP address. They use a nonstandard DNS query operation, and their use is limited to some of the earlier versions of Nslookup, a command-line utility for troubleshooting and testing DNS service.

For Windows 2000 Server, DNS service recognizes and accepts inverse query messages, answering them with a fake inverse query response. For DNS servers running in Windows NT Server 4.0, this support is available by default if the server computer has been updated to Service Pack 4 or later.

The configuration of PTR resource records and reverse lookup zones for identifying hosts by reverse query is strictly an optional part of the DNS standard implementation. You are not required to use reverse lookup zones, although for some networked applications, they are used to perform security checks.

1.7 Windows Clustering

Windows Clustering is a feature of Windows 2000 Advanced Server that provides multiple clustering technologies:

1. **Network Load Balancing Clusters:** Network Load Balancing clusters provide high scalability and availability for TCP/IP based services and applications by combining up to 32 servers running Windows 2000 Advanced Server into a single cluster. The Network Load Balancing service enables Network Load Balancing clusters. Network Load Balancing clusters can also provide load balancing for servers running COM+ applications.
2. **Server Clusters:** Server clusters provide high availability for applications through the failover of resources on servers running Windows 2000 Advanced Server. The Cluster service enables server clusters. You can install Cluster service using the Windows Components wizard, which is part of Windows 2000 Configure Your Server.

Chapter II: Windows NT Networking

2.1 Active Directory Mechanism

Active Directory is the directory service used in Windows 2000 Server and is the foundation of Windows 2000 distributed networks.

Active Directory is the directory service for Windows 2000 Server. It stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory directory service uses a structured data store as the basis for a logical, hierarchical organization of directory information.

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

2.1.1 Introduction

A directory is a hierarchical structure that stores information about objects on the network. A directory service, such as Active Directory, provides the methods for storing directory data and making this data available to network users and administrators. For example, Active Directory stores information about user accounts, such as names, passwords, phone numbers, and so on, and enables other authorized users on the same network to access this information.

2.1.1.1 Directory Services

The Active Directory directory service has the following features:

1. A data store, also known as the directory, which stores information about Active Directory objects. These objects typically include shared resources such as servers, files, printers, and the network user and computer accounts. For more information about the Active Directory data store.

2. A set of rules, the schema that defines the classes of objects and attributes contained in the directory, the constraints and limits on instances of these objects, and the format of their names. For more information about the schema.
3. A global catalog that contains information about every object in the directory. This allows users and administrators to find directory information regardless of which domain in the directory actually contains the data.
4. A query and index mechanism, so that objects and their properties can be published and found by network users or applications.
5. A replication service that distributes directory data across a network. All domain controllers in a domain participate in replication and contain a complete copy of all directory information for their domain. Any change to directory data is replicated to all domain controllers in the domain.
6. Integration with the security subsystem for a secure logon process to a network, as well as access control on both directory data queries and data modifications.
7. To gain the full benefits of Active Directory, the computer accessing the Active Directory over the network must be running the correct client software. To computers not running Active Directory client software, the directory will appear just like a Windows NT directory.

2.2 Domains Overview

A domain defines a security boundary. The directory includes one or more domains, each having its own security policies and trust relationships with other domains. Domains provide several benefits:

1. Security policies and settings (such as administrative rights and access control lists) do not cross from one domain to another.
2. Delegating administrative authority to domains or organizational units eliminates the need for a number of administrators with sweeping administrative authority.
3. Domains help structure your network to better reflect your organization.

4. Each domain stores only the information about the objects located in that domain. By partitioning the directory this way, Active Directory can scale to very large numbers of objects.

Domains are units of replication. All of the domain controllers in a particular domain can receive changes and replicate those changes to all other domain controllers in the domain. A single domain can span multiple physical locations or sites. Using a single domain greatly simplifies administrative overhead.

2.2.1 Domain Trees and Forests

Multiple domains form a forest. Domains can also be combined into hierarchical structures called domain trees.

2.2.1.1 Domain Trees

The first domain in a domain tree is called the root domain. Additional domains in the same domain tree are child domains. A domain immediately above another domain in the same domain tree is referred to as the parent of the child domain.

All domains that have a common root domain are said to form a *contiguous namespace*. This means that the domain name of a child domain is the name of that child domain, added to the name of the parent domain. In this illustration, child.microsoft.com is a child domain of microsoft.com and the parent domain of grandchild.child.microsoft.com. The microsoft.com domain is the parent domain of child.microsoft.com. It is also the root domain of this domain tree.

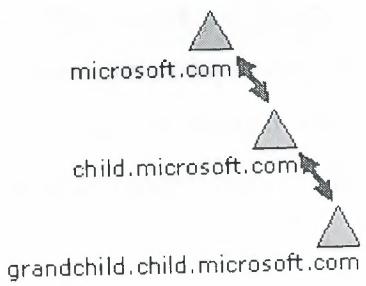

Figure 2.2.1.1: Domain Tree.

Windows 2000 domains in a tree are joined together through two-way, transitive trust relationships. Because these trust relationships are two-way and transitive, a Windows 2000 domain newly created in a domain tree or forest immediately has trust relationships established with every other Windows 2000 domain in the domain tree or forest. These trust relationships allow a single logon process to authenticate a user on all domains in the domain tree or forest. This does not necessarily mean that the authenticated user has rights and permissions in all domains in the domain tree. Because a domain is a security boundary, rights and permissions must be assigned on a per-domain basis.

2.2.1.2 Forests

A forest consists of multiple domain trees. The domain trees in a forest do not form a contiguous namespace. For example, although the two domain trees, microsoft.com and microsoftasia.com may each have a child domain named "support", the DNS names for these child domains would be support.microsoft.com and support.microsoftasia.com. There is no shared namespace.

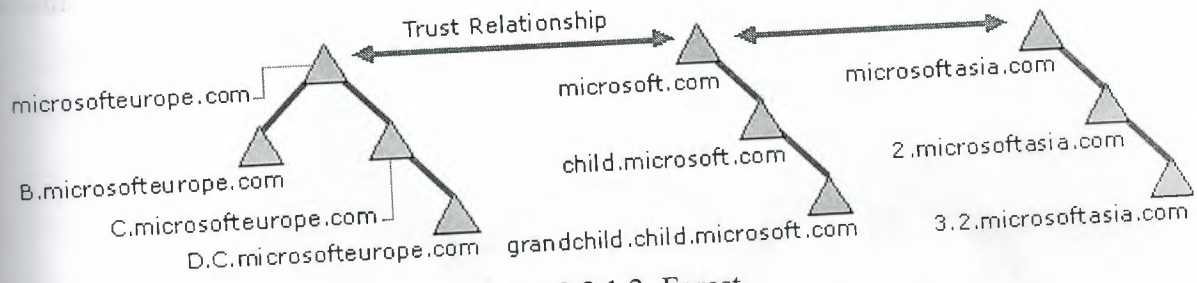

Figure 2.2.1.2: Forest.

However, a forest does have a root domain. The forest root domain is the first domain created in the forest. The root domains of all domain trees in the forest establish transitive trust relationships with the forest root domain. In the illustration, microsoft.com is the forest root domain. The root domains of the other domain trees, microsofteurope.com and microsoftasia.com, have transitive trust relationships with microsoft.com. This is necessary for the purposes of establishing trust across all the domain trees in the forest. All of the Windows 2000 domains in all of the domain trees in a forest share the following traits:

1. Transitive trust relationships between the domains
2. Transitive trust relationships between the domain trees
3. A common schema
4. Common configuration information
5. A common global catalog

Using both domain trees and forests provides you with the flexibility of both contiguous and noncontiguous naming conventions. This can be useful in, for example, companies with independent divisions that must each maintain their own DNS names.

2.2.2 Domain Trusts

A domain trust is a relationship established between two domains that enable users in one domain to be authenticated by a domain controller in another domain. All domain trust relationships have only two domains in the relationship: the trusting domain and the trusted domain.

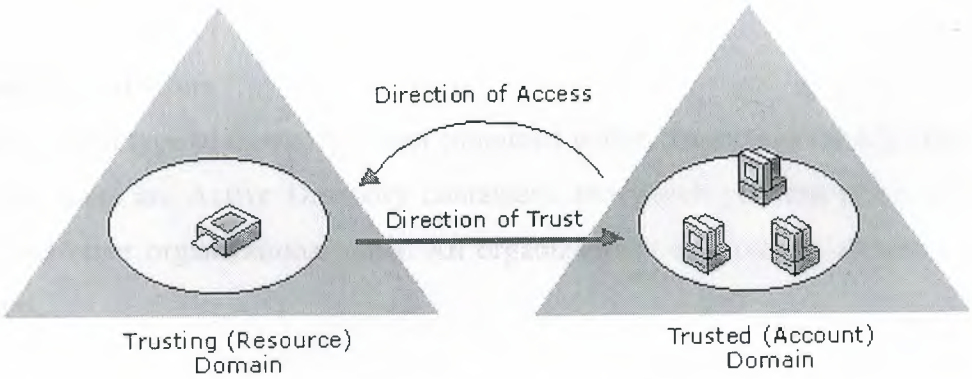

Figure 2.2.2a: Domain Trust.

In the first illustration, trusts are indicated by an arrow (pointing to the trusted domain). In earlier versions of Windows, trusts were limited to the two domains involved in the trust and the trust relationship was one-way. In Windows 2000, all trusts are transitive and two-way. Both domains in a trust relationship automatically trust each other.

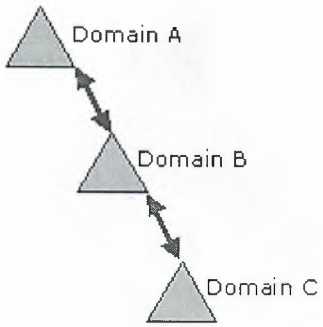

Figure 2.2.2b: Trust Relations on Domain Tree.

As shown in this illustration, this means that if Domain A trusts Domain B and Domain B trusts Domain C, users from Domain C (when granted the proper permissions) can access resources in Domain A.

When a user is authenticated by a domain controller, this does not imply any access to resources in that domain. This is determined solely by the rights and permissions granted to the user account by the domain administrator for the trusting domain.

2.2.3 Organizational Units

A particularly useful type of directory object contained within domains is the organizational unit. Organizational units are Active Directory containers into which you can place users, groups, computers, and other organizational units. An organizational unit cannot contain objects from other domains.

An organizational unit is the smallest scope or unit to which you can assign Group Policy settings or delegate administrative authority. Using organizational units, you can create containers within a domain that represent the hierarchical, logical structures within your organization. This enables you to manage the configuration and use of accounts and resources based on your organizational model.

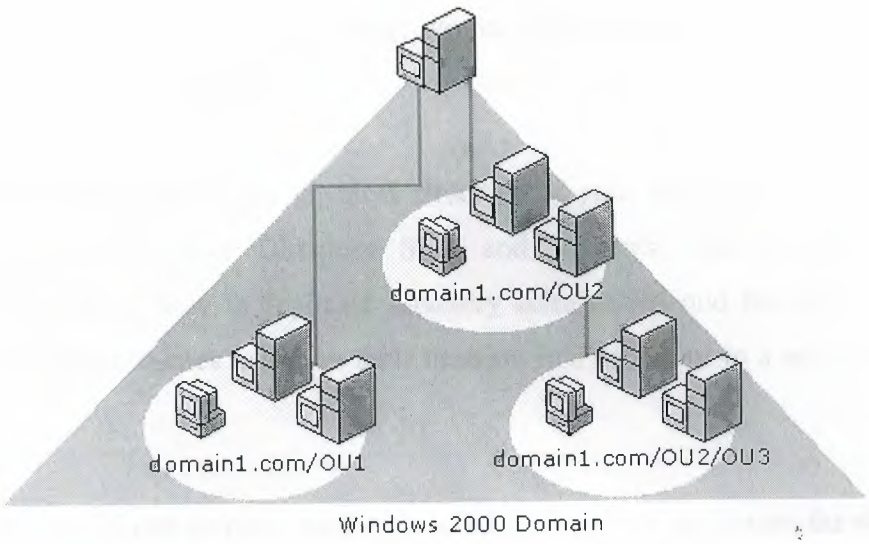

Figure 2.2.3: Organizational Units.

As shown in the illustration, organizational units can contain other organizational units. A hierarchy of containers can be extended as necessary to model your organization's hierarchy within a domain. Using organizational units will help you minimize the number of domains required for your network.

You can use organizational units to create an administrative model that can be scaled to any size. A user can be granted administrative authority for all organizational units in a domain or for a single organizational unit. An administrator of an organizational unit does not need to have administrative authority for any other organizational units in the domain.

2.2.4 Services and Directory Sites

Active Directory uses multimaster replication, enabling any Windows 2000 domain controller in the forest to service requests, including modifications to the directory by users.

If you have a small deployment of well-connected computers, arbitrary selection of a domain controller may not cause problems. However, a deployment that comprises a Wide Area Network (WAN) could be extraordinarily inefficient when, for example, users in Sydney attempt to authenticate to domain controllers in New York using a dial-up connection.

Active Directory Sites and Services can improve the efficiency of directory services for most deployments through the use of sites.

You provide information about the physical structure of your network by publishing sites to Active Directory using Active Directory Sites and Services. Active Directory uses this information to determine how to replicate directory information and handle service requests. Computers are assigned to sites based on their location in a subnet or in a set of well-connected subnets.

Having all computers in one or more well-connected subnets also reinforces the standard that all computers in a site must be well-connected, since computers in the same subnet typically have better connections than an arbitrary selection of computers on a network.

1. **Authentication**. When clients log on using a domain account, the logon mechanism first searches for domain controllers that are in the same site as the client. Attempting to use domain controllers in the client's site first localizes network traffic, increasing the efficiency of the authentication process.
2. **Replication**. Directory information is replicated both within and among sites. Active Directory replicates information within a site more frequently than across sites. This balances the need for up-to-date directory information with the limitations imposed by available network

You customize how Active Directory replicates information using site links to specify how your sites are connected. Active Directory uses the information about how sites are connected to generate Connection objects that provide efficient replication and fault tolerance.

You provide information about the cost of a site link; times when the link is available for use and how often the link should be used. Active Directory uses this information to determine which site link will be used to replicate information. Customizing replication schedules so replication occurs during specific times, such as when network traffic is low, will make replication more efficient.

Ordinarily, all domain controllers are used to exchange information between sites, but you can further control replication behavior by specifying a bridgehead server for inter-site replicated information. Establish a bridgehead server when you have a specific server you want to dedicate for inter-site replication, rather than using any server available. You can also establish a bridgehead server when your deployment uses proxy servers, such as for sending and receiving information through a firewall.

Information such as service bindings and configurations can be made available through the directory, making administration and use of network resources easier and more efficient. Sites help structure and optimize distribution of service information, so the current information is available to clients and distributed efficiently throughout your network.

2.2.5 Groups

Groups are Active Directory or local computer objects that can contain users, contacts, computers, and other groups. Use groups to:

1. Manage user and computer access to shared resources such as Active Directory objects and their properties, network shares, files, directories, printer queues.
2. Filter Group Policy settings.
3. Create e-mail distribution lists.

There are two kinds of groups:

1. Security Groups.
2. Distribution Groups.

Security groups are used to collect users, computers and other groups into manageable units. When assigning permissions for resources (file shares, printers, and so on), administrators should assign those permissions to a security group rather than to individual users. The permissions are assigned once to the group, instead of several times to each individual user.

Each account added to a group receives the rights and permissions defined for that group. Working with groups instead of with individual users helps simplify network maintenance and administration.

Distribution groups can only be used as e-mail distribution lists. They cannot be used to filter Group Policy settings. Distribution groups have no security function. As opposed to groups, organizational units are used to create collections of objects within a single domain, but do not confer membership. The administration of an organizational unit and the objects it contains can be delegated to an individual.

Group Policy objects can be applied to sites, domains or organizational units, but never to groups. A Group Policy object is a collection of settings that affects users or computers. Group membership is used to filter which Group Policy objects will affect the users and computers in the site, domain or organizational unit.

2.2.5.1 Setting User and Group Policy

In order to secure a computer and its resources, you must decide what tasks and actions users or groups of users can perform. The tasks and actions that a user or group of users can perform are determined by the user rights that you assign to them. For example, if a trusted member of the Users group needed to monitor the security log, you could grant the user the "Manage auditing and security log" user right instead of adding the user to a more privileged group, such as the Administrators group. Similarly, you can secure an object, such as a file or folder, by assigning permissions. Some of the most common tasks are assigning user rights on your local computer, assigning user rights throughout your organization, and setting file and folder permissions.

2.2.6 Active Directory Schema

The Active Directory schema is the set of definitions that defines the kinds of objects, and the types of information about those objects, that can be stored in Active Directory. The definitions are themselves stored as objects so that Active Directory can manage the schema objects with the same object management operations used for managing the rest of the objects in the directory.

There are two types of definitions in the schema: attributes and classes. Attributes and classes are also referred to as schema objects or metadata. Attributes are defined separately from classes. Each attribute is defined only once and can be used in multiple classes. For example, the Description attribute is used in many classes, but is defined once in the schema, assuring consistency.

Classes, also referred to as object classes; describe the possible directory objects that can be created. Each class is a collection of attributes. When you create an object, the attributes store the information that describes the object. The User class, for example, is composed of many attributes, including Network Address, Home Directory, and so on. Every object in Active Directory is an instance of an object class.

A set of basic classes and attributes are supplied with Windows 2000 Server. Experienced developers and network administrators can dynamically extend the schema by defining new classes and new attributes for existing classes. Active Directory does not support deletion of schema objects; however, objects can be marked as deactivated, providing many of the benefits of deletion. Extending the schema is an advanced operation with the potential for adverse consequences. The structure and content of the schema is controlled by the domain controller that holds the schema operations master role. A copy of the schema is replicated to all domain controllers in the forest.

The use of this common schema ensures data integrity and consistency throughout the forest. For development and testing purposes, you can also view and modify the Active Directory schema with the Active Directory Schema snap-in, included with the Windows 2000 Administration Tools on the Windows 2000 Server compact disc.

2.3 Active Directory Object Names

Every object in Active Directory is an instance of a class defined in the Active Directory schema.

1. Unique identification of each object (instance of a class) in a directory data store.
2. Backward compatibility with security IDs used in Windows NT 4.0 and earlier.
3. Compatibility with LDAP standards for directory objects names each object in Active Directory can be referenced by several different names. Active Directory creates a relative distinguished name and a canonical name for each object based upon information that was provided when the object was created or modified. Each object can also be referenced by its distinguished name, which is derived from the relative distinguished name of the object and all of its parent container objects. The LDAP relative distinguished name uniquely identifies the object within its parent container. For example, the LDAP relative distinguished name of a computer named my computer is CN=my computer.
4. The canonical name is constructed the same way as the distinguished name, but it is represented using a different notation. The canonical name of the computer in the previous example would be Microsoft.com/MyOrganizationalUnit/mycomputer.

Security principal objects are Active Directory objects that are assigned security identifiers and can be used to log on to the network and can be granted access to domain resources. An administrator needs to provide names for security principal objects (user accounts, computer accounts, and groups) that are unique within a domain.

Consider what occurs when a new user account is added to your directory. You provide a name the user must use to log on to the network, the name of the domain that contains the user account, and other descriptive data, such as first name, last name, telephone number and so on (called attributes). All this information is recorded in the directory.

The names of security principal objects can contain all Unicode characters except the special LDAP characters defined in RFC 2253. This list of special characters includes: a leading space; a trailing space; and any of the following characters: # , + " \ < > ;

From the information provided by the person who creates the security principal object, Active Directory generates a security ID, and a globally unique ID used to identify the security principal. Active Directory also creates an LDAP relative distinguished name, based on the security principal name. An LDAP distinguished name and a canonical name are derived from the relative distinguished name and the names of the domain and container contexts in which the security principal object is created.

If your organization has several domains, it is possible to use the same user name or computer name in different domains. The security ID, globally unique ID, LDAP distinguished name, and canonical name generated by Active Directory will uniquely identify each user, computer, or group in the forest. If the security principal object is renamed or moved to a different domain, the security ID, LDAP relative distinguished name, LDAP distinguished name, and canonical name will change, but the globally unique ID generated by Active Directory will not change.

Security principal objects, such as user accounts, may be renamed, moved, or contained within a nested domain hierarchy. To reduce the effect of renaming, moving, or assigning user account names within a nested domain hierarchy, Active Directory provides a method for simplifying user logon names.

2.4 Active Directory Clients

The Active Directory client is network client software for computers connecting to Active Directory networks. A computer configured with the Active Directory client can log on to the network by locating a domain controller. The client can then fully benefit from the features of Active Directory. Computers with Active Directory clients are:

1. Computers running Windows 2000 Server or Windows 2000 Professional.
2. Computers running Windows 98 or Windows 95 that have add-on Active Directory Installed.

The Active Directory client is provided in a single upgrade pack in a Clients folder on the Windows 2000 Server compact disc.

2.4.1 Locating Domain Name Controller

To logon to an Active Directory network, an Active Directory client must first locate an Active Directory domain controller for their domain. To locate a domain controller for a specified domain, an Active Directory client sends a DNS name query to its configured DNS server(s) with the following characteristics:

1. Query type: SRV (Service locator resource record).
2. Query name: _ldap._tcp.*domain_name*.

The response from the DNS server contains the DNS names of the domain controllers and their IP addresses. From the list of domain controller IP addresses, the client attempts to contact each domain controller to ensure that it is operational. The first domain controller to respond is the domain controller that is used for the logon process.

2.5 Directory Data Store

Active Directory directory service uses a replicated data store. This data store is often simply referred to as the *directory*. The directory contains information about objects such as users, groups, computers, domains, organizational units, and security policies. This information can be published for use by users and administrators.

The directory is stored on domain controllers and can be accessed by network applications or services. A domain can have one or more domain controllers. Each domain controller has a copy of the directory for the domain in which it is located.

Changes made to the directory are replicated from the originating domain controller to other domain controllers in the domain, domain tree, or forest. Because the directory is replicated, and because each domain controller has a copy of the directory, the directory is highly available to users and administrators throughout the domain.

Directory data is stored in the Ntds.dit file on an NTFS partition on the domain controller. Private data is stored securely, and public directory data is stored on a shared system volume

where it can be replicated to other domain controllers in the domain. There are three categories of directory data replicated between domain controllers.

1. **Domain Data:** The domain data contains information about objects within a domain. This is the information typically thought of as directory information such as e-mail contacts, user and computer account attributes, and published resources that are of interest to administrators and users. For example, when a user account is added to your network, a user account object and attribute data are stored in the domain data. When changes to your organization's directory objects occur, such as object creation, deletion, or attribute modification, this data is stored in the domain data.
2. **Configuration Data:** The configuration data describes the topology of the directory. This configuration data includes a list of all domains, trees, and forests, and the locations of the domain controllers and global catalogs.
3. **Schema Data:** The schema is the formal definition of all objects and attribute data that can be stored in the directory. Windows 2000 Server includes a default schema that defines many object types, such as user and computer accounts, groups, domains, organizational units, and security policies. Administrators and programmers can extend the schema by defining new object types and attributes, or by adding new attributes for existing objects. Schema objects are protected by access control lists, ensuring that only authorized users can alter the schema.

2.6 Server Role Management

Windows 2000 Server can operate in any of several roles. You can easily change Windows 2000 Server between the various roles to accommodate the needs of your organization.

2.6.1 Domain Controllers

A domain controller is a computer running Windows 2000 Server that has been configured using the Active Directory Installation wizard. The Active Directory Installation wizard installs and configures components that provide Active Directory directory service to network users and computers. Domain controllers store directory data and manage user-domain interactions, including user logon processes, authentication, and directory searches.

A domain can have one or more domain controllers. A small organization using a single local area network (LAN) may need only one domain with two domain controllers for high availability and fault tolerance. A large company with many network locations will need one or more domain controllers in each location to provide high availability and fault tolerance.

Active Directory supports multimaster replication of directory data between all domain controllers in the domain. Some changes are impractical to perform in multimaster fashion, however, so only one domain controller, called the operations master, accepts requests for such changes. In any Active Directory forest, there are at least five different operations master roles that are assigned to one or more domain controllers.

Windows 2000 Server domain controllers provide an extension of the capabilities and features provided by Windows NT Server 4.0 domain controllers. Windows 2000 Server multimaster replication synchronizes directory data on each domain controller, ensuring consistency of information over time. Multimaster replication is an evolution of the primary and backup domain controller model used in Windows NT Server 4.0, in which only one server, the primary domain controller, had a read and write copy of the directory.

2.6.2 Member Servers

Computers that function as servers within a domain can have one of two roles: domain controller or member server. A member server is a computer that is:

1. Running Windows 2000 Server.
2. A member of a domain.
3. Not a domain controller.

Since it is not a domain controller, a member server does not handle the account logon process, does not participate in Active Directory replication, and does not store domain security policy information. Member servers typically function as the following types of servers:

1. File Servers.
2. Application Servers.
3. Database Servers.
4. Web Servers.
5. Certificate Servers.
6. Firewalls.
7. Remote Access Servers.

These member servers have a common set of security-related features. Member servers adhere to Group Policy settings that are defined for the site, domain, or organizational unit. Resources that are available on a member server are configured for access control. Member server users have user rights assigned to them. Member servers contain a local security account database, the Security Account.

2.6.3 Stand Alone Servers

A stand-alone server is a computer that is running Windows 2000 Server and is not a member of a Windows 2000 domain. If Windows 2000 Server is installed as a member of a workgroup, that server is a stand-alone server.

Stand-alone servers can share resources with other computers on the network, but they do not receive any of the benefits provided by Active Directory. A server within a domain can function in one of two roles: either as a domain controller or a member server.

As the needs of your computing environment change, you might want to change the role of a server. Using the Active Directory Installation wizard, you can promote a member server to a domain controller, or you can demote a domain controller to a member server.

2.7 Benefits of Active Directory Information

Security is fully integrated with Active Directory. Access control can be defined not only on each object in the directory but also on each property of each object. Active Directory provides both the store and the scope of application for security policies.

A security policy can include account information, such as domain-wide password restrictions or rights to particular domain resources. Security policies are implemented through Group Policy settings.

2.7.1 Policy Based Administration

Active Directory service includes both a data store and a logical, hierarchical structure. As a logical structure, it provides a hierarchy of contexts for the application of policy. As a directory, it stores the policies (called Group Policy objects) that are assigned to a particular context. A Group Policy object expresses a set of business rules containing settings that, for the context to which it is applied, can determine:

1. Access to directory objects and domain resources.
2. What domain resources (such as applications) are available to users.
3. How these domain resources are configured for use.

For example, a Group Policy object can determine what applications users see on their computer when they log on, how many users can connect to Microsoft SQL Server when it starts on a server, and what documents or services users can access when they move to different departments or groups. Group Policy objects enable you to manage a small number of policies rather than a large number of users and computers. Active Directory enables you to apply Group Policy settings to the appropriate contexts, whether this is your entire organization or specific units of your organization.

2.7.2 Extensive Functionality

Active Directory is extensible, which means that administrators can add new classes of objects to the schema and can add new attributes to existing classes of objects. For example, you could add a Purchase Authority attribute to the User object and then store each user's purchase authority limit as part of the user's account.

You can add objects and attributes to the directory by using the Active Directory schema or by creating scripts based on Active Directory Service Interfaces (ADSI) or the LDIFDE or CSVDE command-line utilities. For more information

2.7.3 Scalability

Active Directory includes one or more domains, each with one or more domain controllers, enabling you to scale the directory to meet any network requirements. Multiple domains can be combined into a domain tree and multiple domain trees can be combined into a forest.

The directory distributes its schema and configuration information to all domain controllers in the directory. This information is stored in the initial domain controller for a domain and replicated to any additional domain controllers in the domain. When the directory is configured as a single domain, adding domain controllers scales the directory without the administrative overhead involved with additional domains.

Adding domains to the directory enables you to partition the directory for different policy contexts and scale the directory to accommodate a large number of resources and objects.

2.7.4 Information Replication

Replication provides information availability, fault tolerance, load balancing, and performance benefits for the directory. Active Directory uses multimaster replication, enabling you to update the directory at any domain controller, rather than at a single, primary domain controller. The multimaster model has the benefit of greater fault tolerance, since, with multiple domain controllers, replication continues, even if any single domain controller stops working.

Although users may not realize it, due to multimaster replication, they are updating a single copy of the directory. After directory information has been created or modified at a domain controller, the new or changed information is sent to all other domain controllers in the domain, so their directory information is current.

Domain controllers need the latest directory information, but to be efficient, they must limit their updates only to times when there is new or changed directory information. Indiscriminately exchanging directory information among domain controllers can quickly overwhelm any network. Active Directory has been designed to replicate only changed directory information.

With multimaster replication, there is always the potential for the exact same directory change to occur at more than one domain controller. Active Directory has also been designed to track and mediate conflicting changes to the directory, resolving conflicts automatically in nearly all cases.

Deploying multiple domain controllers in one domain provides fault tolerance and load balancing. If one domain controller within a domain slows, stops, or fails, other domain controllers within the same domain can provide necessary directory access, since they contain the same directory data.

2.7.5 DNS Integration

Active Directory uses the Domain Name System (DNS). DNS is an Internet standard service that translates easily readable host names, such as mycomputer.microsoft.com, to numeric IP addresses. This enables identification and connection to processes running on computers on TCP/IP networks.

Domain names for DNS are based on the DNS hierarchical naming structure, which is an inverted tree structure: a single root domain, underneath which can be parent and child domains (branches and leaves). For example, a Windows 2000 domain name such as child.parent.microsoft.com identifies a domain named "child," that is a child domain of the domain named "parent," itself a child of the root domain microsoft.com.

Each computer in a DNS domain is uniquely identified by its DNS fully qualified domain name. The fully qualified domain name of a computer located in the domain child.parent.microsoft.com would be *computername*.child.parent.microsoft.com. Active Directory is integrated with DNS in the following ways:

1. Active Directory and DNS have the same hierarchical structure. Although separate and implemented differently for different purposes, an organization's namespace for DNS and Active Directory have an identical structure. For example, microsoft.com is a DNS domain and an Active Directory domain.
2. DNS zones can be stored in Active Directory. If you are using the Windows 2000 DNS service, primary zone files can be stored in Active Directory for replication to other Active Directory domain controllers.
3. Active Directory clients use DNS to locate domain controllers. To locate a domain controller for a specified domain, Active Directory clients query their configured DNS server for specific resource records.

Since Active Directory is based on standard directory access protocols, such as Lightweight Directory Access Protocol (LDAP) version 3, and the Name Service Provider Interface (NSPI), it can interoperate with other directory services employing these protocols.

LDAP is the directory access protocol used to query and retrieve information from Active Directory. Because it is an industry-standard directory service protocol, programs can be developed using LDAP to share Active Directory information with other directory services that also support LDAP.

The NSPI protocol, which is used by Microsoft Exchange 4.0 and 5.*x* clients, is supported by Active Directory to provide compatibility with the Exchange directory.

2.7.6 Flexible Query

Users and administrators can use the **Search** command on the **Start** menu, **My Network Places** or Active Directory Users and Computers to quickly find an object on the network using object properties. For example, you can find a user by first name, last name, e-mail name, office location, or other properties of that person's user account. Finding information is optimized by use of the global catalog.

2.7.7 Domain Management

Active Directory is an implementation of Internet standard directory and naming protocols. It uses a database engine for transactional support and supports a variety of application programming interface standards. A domain provides several benefits:

1. Organizing Objects: Using organizational units within a domain helps you manage the accounts and resources in the domain.
2. Publishing resources and information about domain objects. Using multiple domains, you can scale the Active Directory directory service to accommodate your administrative and directory publishing requirements.

A domain stores only the information about objects located in that domain, so by creating multiple domains, you are partitioning or segmenting the directory to better serve a disparate user base. Applying a Group Policy object to the domain consolidates resource and security management. A domain defines a scope or unit of policy. A Group Policy object establishes how domain resources can be accessed, configured, and used. These policies are applied only within the domain and not across domains.

Delegating authority eliminates the need for a number of administrators with sweeping administrative authority. Using delegated authority in conjunction with Group Policy objects and group memberships enables you to assign an administrator rights and permissions to manage objects in an entire domain or in one or more organizational units within the domain.

Because a domain is a security boundary, administrative permissions for a domain are limited to the domain by default. For example, an administrator with permissions to set security policies in one domain is not automatically granted authority to set security policies within any other domain in the directory.

To create a domain, you must promote one or more computers running Windows 2000 Server to be domain controllers. A domain controller provides the Active Directory directory service to network users and computers, stores directory data, and manages user-domain interactions, including user logon processes, authentication, and directory searches. Every domain must contain at least one domain controller.

2.7.7.1 Domain Controllers and Forests

Each domain in the directory is identified by a DNS domain name and requires one or more domain controllers. If your network requires more than one domain, you can easily create create multiple domains.

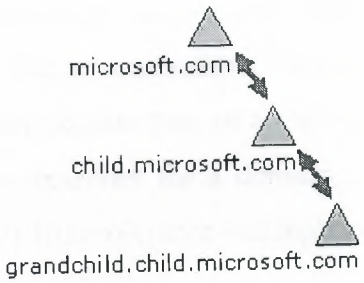

Figure 2.7.7.1a: Domain Tree.

One or more domains that share a common schema and global catalog are referred to as a forest. If multiple domains in the forest have contiguous DNS domain names, as shown in the first illustration, then that structure is referred to as a domain tree.

If, as shown in the second illustration, multiple domains have noncontiguous DNS domain names, then they form separate domain trees within the forest. A forest can contain one or more domain trees. The first domain in a forest is referred to as the forest root domain.

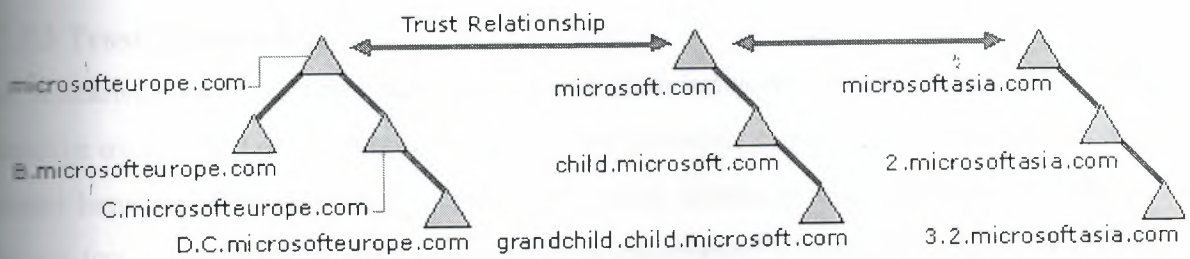

Figure 2.7.7.1b: Forest.

You create a domain by installing the first domain controller for a domain. During installation of the first domain controller, the Active Directory Installation wizard uses information you provide to install the domain controller and create the domain within the existing context (if any) of relationships to other domains and domain controllers. This context may be the first domain in a new forest, the first domain in a new domain tree, or a child domain of an existing domain tree. After you install the first domain controller for a domain, you can install additional domain controllers in an existing domain for fault tolerance and high availability of the directory.

2.7.7.2 Domain Naming

Domains that form a single domain tree share a contiguous namespace (naming hierarchy). Following DNS standards, the fully qualified domain name for a domain that is part of a contiguous namespace is the name of that domain appended to the names of the parent and root domains using the dot (.) character format. For example, a domain with a NetBIOS name of "grandchild" that has a parent domain named parent.microsoft.com would have a fully qualified DNS domain name of grandchild.parent.microsoft.com.

Domain trees associated in a forest share the same Active Directory schema and directory configuration and replication information, but do not share a contiguous DNS domain namespace.

The combination of domain trees and forests provides you with flexible domain naming options. Both contiguous and noncontiguous DNS namespaces can be included in your directory.

2.7.7.3 Trust Relationships

For Windows 2000 computers, account authentication between domains is enabled by two-way, transitive trusts based on the Kerberos V5 security protocol. Trust relationships are automatically created between adjacent domains (parent and child domains) when a domain is created in a domain tree. In a forest, a trust relationship is automatically created between the forest root domain and the root domain of each domain tree added to the forest. Because these trust relationships are transitive, users and computers can be authenticated between any domains in the domain tree or forest.

When upgrading a Windows pre-Windows 2000 domain to Windows 2000, the existing one-way trust relationships between that domain and any other domains are maintained. This includes all trusts with pre-Windows 2000 domains.

If you are installing a new Windows 2000 domain and want trust relationships with any pre-Windows 2000 domains, you must create external trusts with those domains.

2.7.8 Domain and Account Naming

Active Directory domain names are usually the full DNS name of the domain. However, for backward compatibility, each domain also has a pre-Windows 2000 name for use by computers running pre-Windows 2000 operating systems.

The pre-Windows 2000 domain name can be used to log on to a Windows 2000 domain from computers running pre-Windows 2000 operating systems using the *domain name\username* format.

This same format can also be used to log on to a Windows 2000 domain from computers running Windows 2000. Users can also log on to computers running Windows 2000 using the user principal name associated with their user account.

2.7.8.1 User Accounts and Management

In Active Directory, each user account has a user logon name, a pre-Windows 2000 user logon name (security account manager account name), and a user principal name suffix. The administrator enters the user logon name and selects the user principal name suffix when creating the user account. Active Directory suggests a pre-Windows 2000 user logon name using the first 20 bytes of the user logon name.

Administrators can change the pre-Windows 2000 logon name at any time. In Active Directory, each user account has a user principal name based on IETF RFC 822, *Standard for the Format of ARPA Internet Text Messages*. The user principal name is composed of the user logon name and the user principal name suffix joined by the @ sign.

We do not add the @ sign to the user logon name or to the user principal name suffix. Active Directory automatically adds it when it creates the user principal name. A user principal name that contains more than one @ sign is invalid.

The second part of the user principal name, referred to as the user principal name suffix, identifies the domain in which the user account is located. This user principal name suffix can be the DNS domain name, the DNS name of any domain in the forest, or it can be an alternative name created by an administrator and used just for logon purposes. This alternative user principal name suffix does not need to be a valid DNS name.

In Active Directory, the default user principal name suffix is the DNS name of the root domain in the domain tree. In most cases, this is the domain name registered as the enterprise domain on the Internet. Using alternative domain names as the user principal name suffix can provide additional logon security and simplify the names used to log on to another domain in the forest.

For example, if your organization uses a deep domain tree, organized by department and region, domain names can get quite long. The default user principal name suffix for a user in that domain might be *sales.westcoast.microsoft.com*. The logon name for a user in that domain would be *user@sales.westcoast.microsoft.com*. Creating a user principal name suffix of "Microsoft" would allow that same user to log on using the much simpler logon name of *user@microsoft*. You can add or remove user principal name suffixes using Active Directory Domains and Trusts

2.7.8.1.1 Creating User and Group Account

User accounts are used to authenticate, authorize or deny access to resources for, and audit the activity of individual users on your network. A group account is a collection of user accounts that you can use to assign a set of permissions and rights to multiple users simultaneously. A group can also contain contacts, computers, and other groups. You can create user accounts and group accounts in Active Directory to manage domain users. You can also create user accounts and group accounts on a local computer to manage users specific to that computer.

Some of the most common tasks are creating user accounts in Active Directory, creating group accounts in Active Directory, creating user accounts on a local computer, and creating groups on a local computer. You can also use the command line to create user and group accounts in Active Directory or on a local computer.

2.7.8.2 Computer Accounts

Each computer account created in Active Directory has a relative distinguished name, a pre-Windows 2000 computer name (security account manager account name), a primary DNS suffix, a DNS host name and a service principal name. The administrator enters the computer name when creating the computer account. This computer name is used as the LDAP relative distinguished name.

Active Directory suggests the pre-Windows 2000 name using the first 15 bytes of the relative distinguished name. The administrator can change the pre-Windows 2000 name at any time.

The primary DNS suffix defaults to the full DNS name of the domain to which the computer is joined. The DNS host name is built from the first 15 characters of the relative distinguished name and the primary DNS suffix. For example, DNS host name of the computer that is joined to the *mydomain.microsoft.com domain*, and that has the relative distinguished name of *CN=MyComputer1234567890*, would be *mycomputer12345.mydomain.microsoft.com*.

The service principal name is built from the DNS host name. The service principal name is used in the process of mutual authentication between the client and the server hosting a particular service. The client finds a computer account based on the service principal name of the service to which it is trying to connect.

It is possible for administrators to change the way the service principal name is created. This security modification allows a computer to use primary DNS suffixes that are different than the domain to which the computer is joined. The same modification also allows Active Directory to use more than the first 15 bytes of the relative distinguished name when constructing the service principal name.

Computers with these modified computer names will register their names in DNS correctly but an additional procedure is required to enable correct registration of the DNS host name (dNSHostName) and service principal Name (servicePrincipalName) attributes of the computer object in Active Directory.

2.8 Domain Trusts

A domain trust is a relationship established between domains that enable users in one domain to be authenticated by a domain controller in the other domain. The authentication requests follow a *trust path*.

2.8.1 Trust Paths

A trust path is the series of trust relationships that authentication requests must follow between domains. Before a user can access a resource in another domain, Windows 2000 security must determine whether the trusting domain (the domain containing the resource the user is trying to access) has a trust relationship with the trusted domain (the user's logon domain).

To determine this, the Windows 2000 security system computes the trust path between a domain controller in the trusting domain and a domain controller in the trusted domain. In the illustration, trust paths are indicated by arrows showing the direction of the trust:

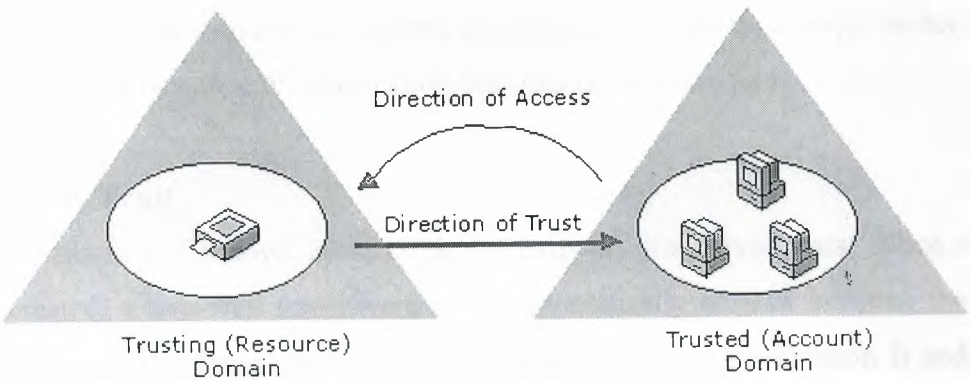

Figure 2.8.1: Trust Paths.

All domain trust relationships have only two domains in the relationship: the trusting domain and the trusted domain. A domain trust relationship is characterized by whether it is:

1. Two-way.
2. One-way.
3. Transitive.
4. No transitive.

2.8.1.1 One Way Trust

A one-way trust is a single trust relationship, where Domain A trusts Domain B. All one-way relationships are no transitive and all no transitive trusts are one-way. Authentication requests can only be passed from the trusting domain to the trusted domain. This means that if Domain A has a one-way trust with Domain B and Domain B has a one-way trust with Domain C; Domain A does not have a trust relationship with Domain C. A Windows 2000 domain can establish a one-way trust with:

1. Windows 2000 domains in a different forest.
2. Windows NT 4.0 domains.
3. MIT Kerberos V5 realms.

Since all Windows 2000 domains in a forest are linked by transitive trust, it is not possible to create one-way trusts between Windows 2000 domains in the same forest.

2.8.1.2 Two Way Trust

All domain trusts in a Windows 2000 forest are two-way transitive trusts. When a new child domain is created, a two-way transitive trust is automatically created between the new child domain and the parent domain. In a two-way trust, Domain A trusts Domain B and Domain B trusts Domain A. This means that authentication requests can be passed between the two domains in both directions.

To create a no transitive two-way trust, you must create two one-way trusts between the domains involved.

2.8.2 Transitive Trust

All domain trusts in a Windows 2000 forest are transitive. Transitive trusts are always two-way: Both domains in the relationship trust each other. A transitive trust is not bounded by the two domains in the trust relationship.

Each time you create a new child domain, a two-way transitive trust relationship is created implicitly (automatically) between the parent and new child domain. In this way, transitive trust relationships flow upward through the domain tree as it is formed, creating transitive trust between all domains in the domain tree. Each time you create a new domain tree in a forest, a two-way transitive trust relationship is created between the forest root domain and the new domain (the root of the new domain tree).

If no child domains are added to the new domain, the trust path is between this new root domain and the forest root domain. If child domains are added to the new domain (making it a domain tree), trust flows upward through the domain tree to the domain tree's root domain, extending the initial trust path created between the domain root and the forest root domain.

Whether the new domain added to the forest is a single root domain (having no child domains) or a domain tree, the trust paths extend through the forest root domain to any other root domains in the forest. In this way, transitive trust relationships flow through all domains in the forest.

Authentication requests follow these trust paths, so accounts from any domain in the forest can be authenticated at any other domain in the forest. With a single logon process, those accounts having the proper permissions can potentially access resources on any domain in the forest.

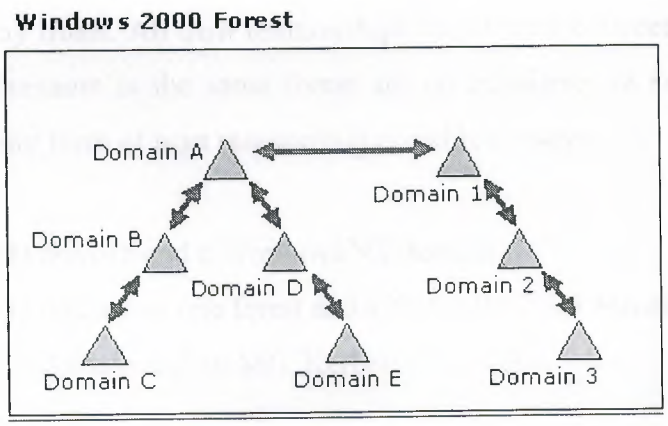

Figure 2.8.2: Transitive Trust Relationship.

Since Domain 1 has a transitive trust relationship with Domain 2 and Domain 2 has a transitive trust relationship with Domain 3, users in Domain 3 (when granted the proper permissions) can access resources in Domain 1. Because Domain 1 has a transitive trust relationship with Domain A, and the other domains in Domain A's domain tree have transitive trust relationships with Domain A, users in Domain B (when granted the proper permissions) can access resources in Domain 3.

You can also explicitly (manually) create transitive trusts between Windows 2000 domains in the same domain tree or forest. These shortcut trust relations can be used to shorten the trust path in large and complex domain trees or forests.

2.8.3 Non transitive Trust

A no transitive trust is bounded by the two domains in the trust relationship and does not flow to any other domains in the forest. In most cases, you must explicitly create no transitive trusts.

All trust relationships between Windows 2000 domains and Windows NT domains are no transitive. When upgrading from Windows NT to Windows 2000, all existing Windows NT trusts are preserved intact. In a mixed-mode environment, all Windows NT trusts are no transitive.

No transitive trusts are one-way by default, although you can also create a two-way relationship by creating two one-way trusts. All trust relationships established between domains that are not both Windows 2000 domains in the same forest are no transitive. In summary, no transitive domain trusts are the only form of trust relationship possible between:

1. A Windows 2000 domain and a Windows NT domain.
2. A Windows 2000 domain in one forest and a Windows 2000 domain in another forest.
3. A Windows 2000 domain and an MIT Kerberos V5 realm.

2.8.4 Trust Protocols

Windows 2000 authenticates users and applications using one of two protocols: Kerberos V5 or NTLM. The Kerberos V5 protocol is the default protocol for computers running Windows 2000 and computers with Windows 2000 client software installed.

If any computer involved in a transaction does not support Kerberos V5, the NTLM protocol will be used. With the Kerberos V5 protocol, the client requests a ticket from a domain controller in its account domain to the server in the trusting domain. This ticket is issued by an intermediary trusted by the client and the server. The client presents this trusted ticket to the server in the trusting domain for authentication.

When a client tries to access resources on a server in another domain using NTLM authentication, the server containing the resource must contact a domain controller in the client's account domain to verify the account credentials.

2.8.5 Explicit Domain Trusts

Explicit trusts are trust relationships that you create yourself, as opposed to trusts created automatically during installation of a domain controller. You create and manage explicit trusts using Active Directory Domains and Trusts. There are two kinds of explicit trusts: external trusts and shortcut trusts. External trusts enable user authentication to a domain outside of a forest. Shortcut trusts shorten the trust path in a complex forest.

2.8.5.1 External Trust

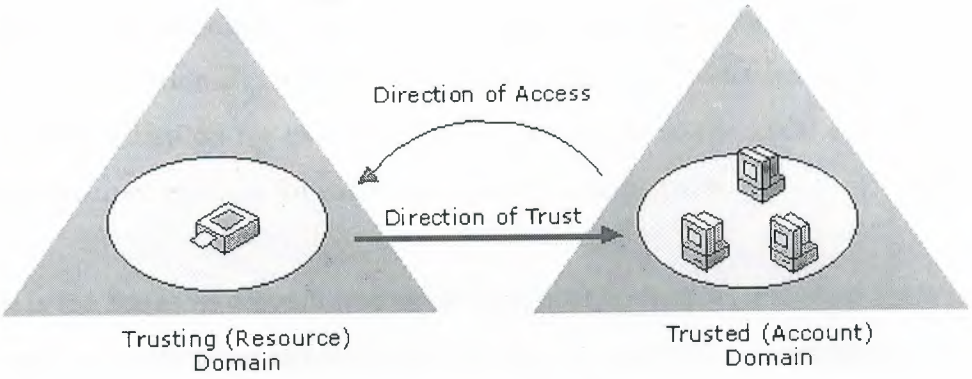

Figure 2.8.5.1: External Trust Relationship.

External trusts create trust relationships to domains outside the forest. The benefit of creating external trusts is to enable user authentication to a domain not encompassed by the trust paths of a forest. All external trusts are one-way non-transitive trusts, as shown in the figure. You can combine two one-way trusts to create a two-way trust relationship.

In mixed-mode domains, external trusts should always be deleted from a Windows 2000 domain controller. External trusts to Windows NT 4.0 or 3.51 domains can be deleted by authorized administrators on the Windows NT 4.0 or 3.51 domain controllers. However, only the trusted side of the relationship can be deleted on Windows NT 4.0 or 3.51 domain controllers.

The trusting side of the relationship (created in the Windows 2000 domain) is not deleted, and although it will not be operational, the trust will continue to display in Active Directory Domains and Trusts. To remove the trust completely, you will need to delete the trust from a Windows 2000 domain controller in the trusting domain.

If an external trust is inadvertently deleted from a Windows NT 4.0 or 3.51 domain controller, you will need to recreate the trust from any Windows 2000 domain controller in the trusting domain.

2.8.5.2 Shortcut Trusts

Before an account can be granted access to resources by a domain controller of another domain, Windows 2000 must determine whether the domain containing the desired resources (target domain) has a trust relationship with the domain in which the account is located (source domain). To make this determination for two domains in a forest, Windows 2000 computes a trust path between the domain controllers for these source and target domains.

A trust path is the series of domain trust relationships that must be traversed by Windows 2000 security to pass authentication requests between any two domains. Computing and traversing a trust path between domains trees in a complex forest can take time, which can be reduced with shortcut trusts.

Shortcut trusts are two-way transitive trusts that enable you to shorten the path in a complex forest. You explicitly create shortcut trusts between Windows 2000 domains in the same forest. A shortcut trust is a performance optimization that shortens the trust path for Windows 2000 security to take for authentication purposes. The most effective use of shortcut trusts is between two domain trees in a forest.

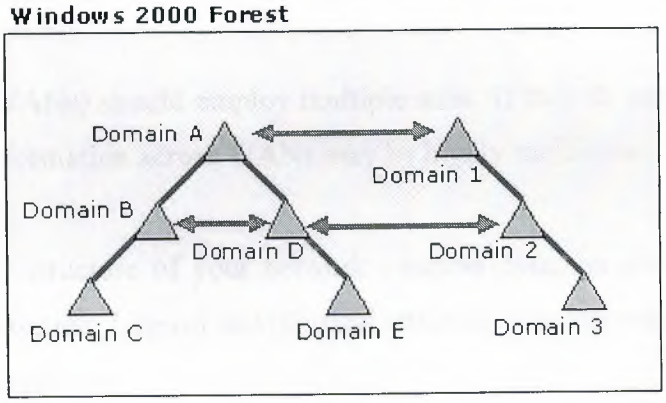

Figure 2.8.5.2: Shortcut Trust Relationship.

As shown in the illustration, you can create a shortcut trust between mid-level domains in two domain trees to shorten the trust path between two Windows 2000 domains in a forest and optimize the Windows 2000 authentication process.

2.8.5.3 Creating Explicit Trusts

To create an explicit trust, you must know the domain names and a user account with permission to create trusts in each domain. Each trust is assigned a password that must be known to the administrators of both domains in the relationship.

2.9 Site and Domain Relation

Figure 2.9a: Site and Domain Relation.

For convenience, think of sites as being defined by a set of computers in one or more IP subnets. This works well because, for efficient exchange of directory information, computers in a site need to be well-connected, a typical characteristic of computers within a subnet. If a site comprises multiple subnets, those subnets too must be well-connected for the same reason.

Wide area networks (WANs) should employ multiple sites. If they do not, servicing requests or replicating directory information across WANs may be highly inefficient.

Sites map the physical structure of your network whereas domains generally map the logical structure of your organization. Logical and physical structure is independent of each other, which has the following consequences:

1. There is no necessary correlation between your network's physical structure and its domain structure.
2. Active Directory allows multiple domains in a single site, as well as multiple sites in a single domain.

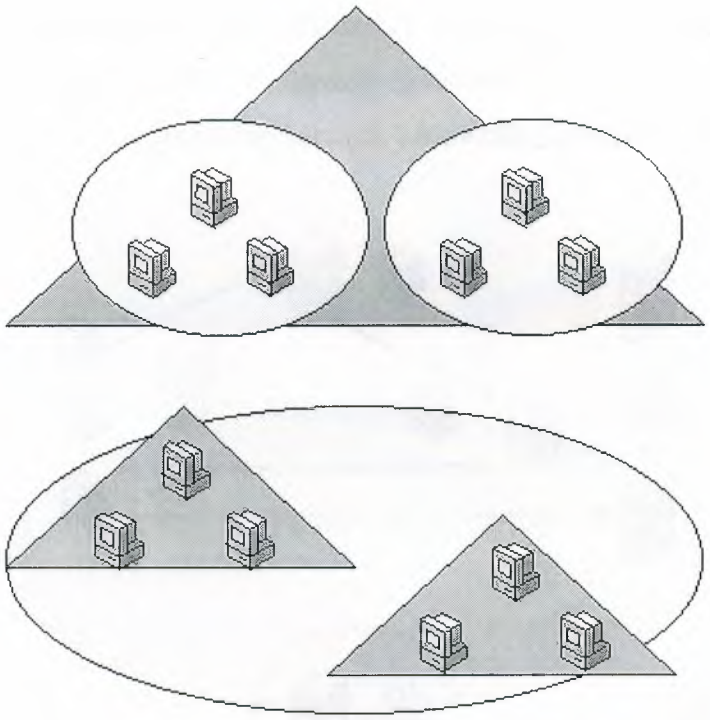

Figure 2.9b: Site and Domain Relation.

2.9.1 Site Management

Active Directory Sites and Services allow you to specify site information. Active Directory uses this information to determine how best to use available network resources. This makes the following types of operations more efficient:

When a client requests a service from a domain controller, it directs the request to a domain controller in the same site, if one is available. Selecting a domain controller that is well-connected to the client that placed the request makes handling the request more efficient.

Sites streamline replication of directory information. Directory schema and configuration information is distributed throughout the forest and domain data is distributed among all domain controllers in the domain. By strategically reducing replication, the strain on your network can be similarly reduced. Active Directory replicates directory information within a site more frequently than among sites.

By this way, the best connected domain controllers, those most likely to need particular directory information, receive replications first. The domain controllers in other sites receive all changes to the directory, but less frequently, reducing network bandwidth consumption.

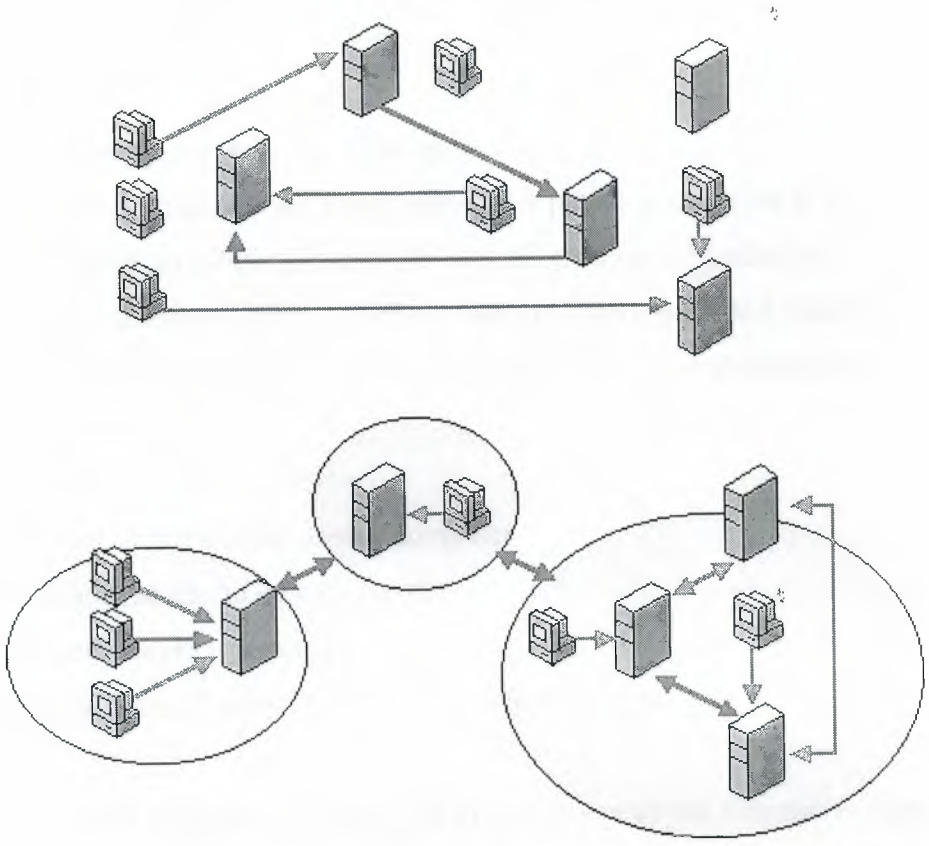

Figure 2.9.1: Site Management.

If a deployment is not organized into sites, information exchange among domain controllers and clients can be chaotic. Sites improve the efficiency of network usage. Site membership is determined differently for domain controllers and clients.

A client determines what site it is in when it is turned on, so its site location will often be dynamically updated. A domain controller's site location is established by which site its Server object belongs to in the directory, so its site location will be consistent unless the domain controller's Server object is intentionally moved to a different site. If a domain controller or client has an address that is not included in any site, then the client or domain controller is contained within the initial site created (Default-First-Site).

All activity is then handled as though the client or domain controller activity is a member of Default-First-Site, regardless of actual IP address or subnet location.

Therefore, all sites will always have a domain controller associated, since the nearest domain controller associates itself with a site that has no domain controller (unless the site Default-First-Site is deleted).

2.10 Active Directory and User Account Management

Active Directory user and computer accounts represent a physical entity such as a computer or person. User accounts and computer accounts (as well as groups) are called security principals. Security principals are directory objects that are automatically assigned security identifiers. Objects with security identifiers can log on to the network and access domain resources. A user or computer account is used to:

1. Authenticate the identity of the user or computer.
2. Authorize or deny access to domain resources.
3. Administer other security principals.
4. Audit actions performed using the user or computer account.

For example, the user and computer accounts that are members of the Enterprise Admins group are automatically granted permission to log on at all of the domain controllers in the forest. User and computer accounts are added, disabled, reset, and deleted using Active Directory Users and Computers.

When a trust is established between a Windows 2000 domain in a particular forest and a Windows 2000 domain outside of that forest, security principals from the external domain can be granted access to resources in the forest. Active Directory creates a "foreign security principal" object to represent each security principal from the trusted external domain. These foreign security principals can become members of domain local groups, which can have members from domains outside of the forest.

Directory objects for foreign security principals are created by Active Directory and should not be manually modified. You can view foreign security principal objects from Active Directory Users and Computers by enabling Advanced Features.

When a trust is established between a Windows 2000 domain in a particular forest and a Windows 2000 domain outside of that forest, security principals from the external domain can be granted access to resources in the internal domain.

Active Directory creates a "foreign security principal" object in the internal domain to represent each security principal from the trusted external domain. These foreign security principals can become members of domain local groups in the internal domain. (Domain local groups can have members from domains outside of the forest.)

2.10.1 Account Management

An Active Directory user account enables a user to log on to computers and domains with an identity that can be authenticated and authorized for access to domain resources. Each user who logs on to the network should have his or her own unique user account and password. User accounts can also be used as service accounts for some applications.

Windows 2000 provides predefined user accounts that you can use to log on to a computer running Windows 2000. These predefined accounts are:

1. Administrator Account.
2. User Account.
3. Guest account.

Predefined accounts are default user accounts designed to let users log on to a local computer and access resources on the local computer. These are designed primarily for initial logon and configuration of a local computer. Each predefined account has a different combination of rights and permissions. The Administrator account has the most extensive rights and permissions while the Guest account has limited rights and permissions.

If predefined account rights and permissions are not modified or disabled by a network administrator, they could be used by any user or service to log on to a network using the Administrator or Guest identity. To obtain the security of user authentication and authorization, create an individual user account for each user who will participate on your network by using Active Directory Users and Computers. Each user account (including the Administrator and Guest account) can then be added to Window 2000 groups to control the rights and permissions assigned to the account.

Using accounts and groups that are appropriate for your network ensures that users logging on to a network can be identified and can access only the permitted resources.

2.10.2 User Account Options

Each Active Directory user account has a number of security related options that determine how someone logging on with that particular user account is authenticated on the network. Several of these options are specific to passwords:

1. User must change password at next logon.
2. User cannot change password.
3. Password never expires.
4. Store passwords using reversible encryption.

These options are self-explanatory except for **Store passwords using reversible encryption**. If you have users logging on to your Windows 2000 network from Apple computers, select this option for those user accounts.

Select the **Account disabled** option to prevent users from logging on with the selected account. Many administrators use disabled accounts as templates for common user accounts. You can use the remaining options to configure security-specific information for Active Directory user accounts:

1. Smart card required for interactive logon.
2. Account is trusted for delegation.
3. Account is sensitive and cannot be delegated.
4. Use DES encryption types for this account.
5. Don't require Kerberos preauthentication.

Select the **Smart card required for interactive logon** option to securely store public and private keys, passwords, and other types of personal information for this user account. There must be a smart card reader attached to the user's computer, and they must have a personal identification number (PIN) to be able to log on to the network.

Select the **Account is trusted for delegation** option to give a user the ability to assign responsibility for management and administration of a portion of the domain namespace to another user, group, or organization.

Select the **Account is sensitive and cannot be delegated** option if this account cannot be assigned for delegation by another account. Select the **don't require Kerberos preauthentication** option if the account uses another implementation of the Kerberos protocol. Not all implementations or deployments of the Kerberos protocol use this feature.

The Kerberos Key Distribution Center uses ticket-granting tickets for obtaining network authentication in a domain. The time at which the Key Distribution Center issues a ticket-granting ticket is important to the Kerberos protocol. Windows 2000 uses other mechanisms to synchronize time, so using the Kerberos preauthentication option works well.

Select the **Use DES encryption types for this account** option if you need the Data Encryption Standard (DES). DES supports multiple levels of encryption, including MPPE Standard (40-bit), MPPE Standard (56-bit), MPPE Strong (128-bit), IPSec DES (40-bit), IPSec 56-bit DES, and IPSec Triple DES (3DES).

2.10.3 Computer Accounts

Every computer running Windows 2000 or Windows NT that joins a domain has a computer account. Similar to user accounts, computer accounts provide a means for authenticating and auditing the computer's access to the network, and access to domain resources. Each computer connected to the network should have its own unique computer account. Computer accounts are also created using Active Directory Users and Computers.

Computers running Windows 98 and Windows 95 do not have the advanced security features of those running Windows 2000 and Windows NT, and cannot be assigned computer accounts in Windows 2000 domains. However, you can log on to a network and use Windows 98 and Windows 95 computers in Active Directory domains.

2.11 Group Policy Management

Group Policy settings affect computer or user accounts and can be applied to sites, domains, or organizational units. It can be used to configure security options, manage applications, manage desktop appearance, assign scripts, and redirect folders from local computers to network locations. Examples:

1. Set the minimum password length and the maximum length of time that a password will remain valid. This can be configured for an entire domain.
2. Administrators can automatically install an application on every computer in a particular domain or on all computers assigned to a particular group in a particular site. For example, you could automatically install Microsoft Outlook on every computer in the domain and automatically install Microsoft Excel only on those computers belonging to the Accounting group in a particular site.
3. Unique logon and logoff scripts can be assigned to the user accounts in each organizational unit.
4. If members of a particular group often use different computers, administrators can install the necessary applications on each of those computers.
5. Any user's My Documents folder can be redirected to a network location. Users and then gain access to their documents from any computer on the network.

2.12 DNS Integration

While Active Directory is integrated with DNS and shares the same namespace structure, it is important to note the difference between them:

1. DNS is a name resolution service. DNS clients send DNS name queries to their configured DNS server. The DNS server receives the name query and either resolves the name query through locally stored files or consults another DNS server for resolution. DNS does not require Active Directory to function.
2. Active Directory is a directory service Active Directory provides an information repository and services to make information available to users and applications. Active Directory clients send queries to Active Directory servers using the Lightweight Directory Access Protocol (LDAP). In order to locate an Active Directory server, an Active Directory client queries DNS. Active Directory requires DNS to function.

Active Directory uses DNS as a locator service, resolving Active Directory domain, site, and service names to an IP address. To log on to an Active Directory domain, an Active Directory client queries their configured DNS server for the IP address of the LDAP service running on a domain controller for a specified domain.

2.12.1 DNS Server and Active Directory Requirements

In order for Active Directory to function properly, DNS servers must provide support for Service Location (SRV) resource records described in RFC 2052, *A DNS RR for specifying the location of services (DNS SRV)*. SRV resource records map the name of a service to the name of a server offering that service.

Active Directory clients and domain controllers use SRV records to determine the IP addresses of domain controllers. Although not a technical requirement of Active Directory, it is highly recommended that DNS servers provide support for DNS dynamic updates described in RFC 2136, *Observations on the use of Components of the Class an Address Space within the Internet.*

The Windows 2000 DNS service provides support for both SRV records and dynamic updates. If a non-Windows 2000 DNS server is being used, verify that it at least supports the SRV resource record. If not, it must be upgraded to a version that does support the use of the SRV resource record. For example, Windows NT Server 4.0 DNS servers must be upgraded to Service Pack 4 or later to support SRV resource records. A DNS server that supports SRV records but does not support dynamic update must be updated with the contents of the Netlogon.dns file created by the Active Directory Installation wizard while promoting a Windows 2000 Server to a domain controller. The Netlogon.dns file is described in the following section.

2.12.2 DNS and Active Directory Installation

By default, the Active Directory Installation wizard attempts to locate an authoritative DNS server for the domain being configured from its list of configured DNS servers that will accept a dynamic update of an SRV resource record. If found, all the appropriate records for the domain controller are automatically registered with the DNS server after the domain controller is restarted.

If a DNS server that can accept dynamic updates is not found, either because the DNS server does not support dynamic updates or dynamic updates are not enabled for the domain, the following steps are taken to ensure that the installation process is completed with the necessary registration of the SRV resource records:

1. The DNS service is installed on the domain controller and is automatically configured with a zone based on the Active Directory domain. For example, if the Active Directory domain that you chose for your first domain in the forest was example.microsoft.com, a zone rooted at the DNS domain name of example.microsoft.com is added and configured to use the DNS service on the new domain controller.
2. A text file containing the appropriate DNS resource records for the domain controller is created. The file called Netlogon.dns is created in the %systemroot%\System32\config folder and contains all the records needed to register the resource records of the domain controller. Netlogon.dns is used by the Windows 2000 NetLogon service and to support Active Directory for non-Windows 2000 DNS servers.

If you are using a DNS server that supports the SRV resource record but does not support dynamic updates (such as a UNIX-based DNS server or a Windows NT Server 4.0 DNS server), you can import the records in Netlogon.dns into the appropriate primary zone file to manually configure the primary zone on that server to support Active Directory.

2.13 Group Types

There are two types of groups in Windows 2000:

1. Security groups.
2. Distribution groups.

Security groups are listed in discretionary access control lists (DACLs) that define permissions on resources and objects. Security groups can also be used as an e-mail entity. Sending an e-mail message to the group sends the message to all the members of the group.

Distribution groups are not security-enabled. They cannot be listed in DACLs. Distribution groups can be used only with e-mail applications (such as Exchange), to send e-mail to collections of users. If you do not need a group for security purposes, create a distribution group instead of a security group.

A group can be converted from a security group to a distribution group, and vice versa, at any time, but only if the domain is in native-mode. No groups can be converted while a domain is in mixed-mode.

2.13.1 Group Scopes

Each security and distribution group has a scope that identifies the extent to which the group is applied in the domain tree or forest. There are three different scopes: universal, global, and domain local.

1. Groups with universal scope can have as their members groups and accounts from any Windows 2000 domain in the domain tree or forest and can be granted permissions in any domain in the domain tree or forest. Groups with universal scope are referred to as universal groups.
2. Groups with global scope can have as their members groups and accounts only from the domain in which the group is defined and can be granted permissions in any domain in the forest. Groups with a global scope are referred to as global groups.
3. Groups with domain local scope can have as their members groups and accounts from a Windows 2000 or Windows NT domain and can be used to grant permissions only within a domain. Groups with a domain local scope are referred to as domain local groups.

If you have multiple forests, users defined in only one forest cannot be placed into groups defined in another forest, and groups defined in only one forest cannot be assigned permissions in another forest.

2.13.1.1 Changing Group Scopes

When creating a new group, by default, the new group is configured as a security group with global scope regardless of the current domain mode. Although changing a group scope is not allowed in mixed-mode domains, the following conversions are allowed in native-mode domains:

1. **Global to Universal**. However, this is only allowed if the group is not a member of another group having global scope.
2. **Domain local to universal**. However, the group being converted cannot have as its member another group having domain local scope.

Several default groups are installed in the built-in and Users folders of the Active Directory Users and Computers console when you install a domain controller. These groups are security groups and represent common sets of rights and permissions that you can use to grant certain roles, rights, and permissions to the accounts and groups that you place into the default groups.

Default groups with domain local scope are located in the Built-in folder. Predefined groups with global scope are located in the Users folder. You can move the Built-in and predefined groups to other group or organizational unit folders within the domain, but you cannot move them to other domains.

2.13.2 Built-in Groups

The default groups placed in the Built-in folder for Active Directory Users and Computers are:

1. Administrators.
2. Backup Operators.
3. Guests.
4. Print Operators.
5. Replicator.
6. Server Operators.
7. Users.

These built-in groups have domain local scope and are primarily used to assign default sets of permissions to users who will have some administrative control in that domain. For example, the Administrators group in a domain has a broad set of administrative authority over all accounts and resources in the domain.

2.13.3 Predefined Groups

The predefined groups placed in the Users folder for Active Directory Users and Computers are:

1. Group Name.
2. Cert Publishers.
3. Domain Admins.
4. Domain Computers.
5. Domain Controllers.
6. Domain Guests.
7. Domain Users.
8. Enterprise Admins.

9. Group Policy Admins.
10. Schema Admins.

You can use these groups with global scope to collect the various types of user accounts in that domain (regular users, administrators, and guests) into groups. These groups can then be placed in groups with domain local scope in that domain and others.

By default, any user account you create in a domain is automatically added to the Domain Users group and any computer account you create is automatically added to the Domain Computers group. You can use the Domain Users and Domain Computers groups to represent all the accounts created in the domain.

For example, if you want all the users in this domain to have access to a printer, you can assign permissions for the printer to the Domain Users group (or put the Domain Users group into a domain local group that has permissions for the printer). By default, the Domain Users group in a domain is a member of the Users group in the same domain.

The Domain Admins group can represent the users who have broad administrative rights in a domain. Windows 2000 Server does not place any accounts in this group automatically, but if you want an account to have sweeping administrative rights in a domain (and possibly other domains), you can put that account into Domain Admins.
Because Windows 2000 Server supports delegation of authority, you should not have to grant these broad administrative rights to many users.

By default, the Domain Admins group in a domain is a member of the Administrators group in the same domain. By default, the Domain Guests group is a member of the Guests group in the same domain, and automatically has as its member the domain's default Guest user account.

In addition to the groups in the Built-in and Users folder, Windows 2000 Server includes several special identities. For convenience, these identities are generally referred to as groups. These special groups do not have specific memberships that you can modify, but they can represent different users at different times, depending on the circumstances. The three special groups are:

1. Everyone: Represents all current network users, including guests and users from other domains. Whenever a user logs on to the network, they are automatically added to the everyone group.
2. Network: Represents users currently accessing a given resource over the network (as opposed to users who access a resource by logging on locally at the computer where the resource is located). Whenever a user accesses a given resource over the network, they are automatically added to the Network group.
3. Interactive: Represents all users currently logged on to a particular computer and accessing a given resource located on that computer (as opposed to users who access the resource over the network).

Although the special identities can be assigned rights and permission to resources, you cannot modify or view the memberships of these special identities. You do not see them when you administer groups and cannot place the special identities into groups. Group scopes do not apply to special identities. Users are automatically assigned to these special identities whenever they log on or access a particular resource.

2.13.4 Groups and Windows 2000 Stand Alone Servers

Some group features, such as universal groups, group nesting, and the distinction between security groups and distribution groups, are available only on Active Directory domain controllers and member servers. Group accounts on Windows 2000 Professional and Windows 2000 Server stand-alone servers work the same way as in Windows NT 4.0:

1. Only local groups are available to be created locally on the computer.
2. A local group created on one of these computers can be assigned permissions only on that one computer.

A Windows 2000 Professional computer that joins a Windows 2000 domain gets added benefits from the domain. Global groups and universal groups from that domain, as well as global groups and universal groups from all domains in the forest, can be displayed. You can assign permissions for the local computer to these groups or place them in the local computer groups.

2.13.5 Nested Groups

Using nesting, you can add a group as a member of another group. You can nest groups to consolidate group management by increasing the affected member accounts and to reduce replication traffic caused by replication of group membership changes.

Your nesting options depend on whether the domain is in native mode or mixed-mode. Groups in native-mode domains or distribution groups in mixed-mode domains have their membership determined as follows:

1. Groups with universal scope can have as their members: accounts, computer accounts, other groups with universal scope, and groups with global scope from any domain.
2. Groups with global scope can have as their members: accounts from the same domain and other groups with global scope from the same domain.
3. Groups with domain local scope can have as their members: accounts, groups with universal scope, and groups with global scope, all from any domain. They can also have as members other groups with domain local scope from within the same domain. Security groups in a mixed-mode domain are restricted to the following types of membership:
4. Groups with global scope can have as their members only accounts.
5. Groups with domain local scope can have as their members other groups with global scope and accounts.

Security groups with universal scope cannot be created in mixed-mode domains because universal scope is supported only in Windows 2000 native-mode domains

2.13.6 Performance Measures

When a user logs on to a Windows 2000 network, the Windows 2000 domain controller determines which groups the user belongs to. Windows 2000 creates a security token and assigns it to the user. The security token lists the user account ID and the security IDs of all the security groups the user belongs to. Group membership can impact network performance through:

1. Replication of groups with universal scope.
2. Network bandwidth.

Building the security token takes time, so the more security groups that a user belongs to, the longer it will take to build that user's security token and the longer it will take that user to log on to the network. The significance of this effect will vary depending upon network bandwidth as well as the configuration of the domain controller that handles the log on process.

Sometimes you may want to create a group for e-mail purposes only, with no intention of using that group to assign rights and permissions to its members. To improve logon performance, create such groups as distribution groups instead of security groups. Because distribution groups are ignored when Windows 2000 builds the user security token during the logon process, this reduces both the size of the token, and the time it takes to build it.

2.13.7 Universal Group Replication

Changes to the data stored in the global catalog are replicated to every global catalog in the forest. Groups having universal scope *and their members* are listed in the global catalog. Whenever one member of a group with universal scope changes, the entire group membership must be replicated to all global catalogs in the domain tree or forest.

Groups having global or domain local scope are also listed in the global catalog, but their members are not. This reduces the size of the global catalog, and dramatically reduces the replication traffic needed to keep the global catalog up to date. You can improve network performance by using groups with global or domain local scope for directory objects that will change frequently.

2.14 Network Bandwidth

Each user's security token is sent to every computer that the user accesses so the target computer can determine whether the user has any rights or permissions at that computer by comparing all the security IDs contained in the token against the permissions listed for any resources at that computer. The target computer also checks whether any of the security IDs in the token belong to any local groups at the target computer.

The more groups a user belongs to, the larger their security tokens will be. If your network has a large number of users, the effect of these large security tokens on your network bandwidth and domain controller processing capability can be significant.

For example, suppose a particular domain contains 500 file shares, each with a corresponding assignment for a group with domain local scope used to grant read access. If most users have read access to most shares, then most employees will have about 500 group security IDs added to their token. This can take a significant amount of time and add considerable data traffic to the network.

2.15 Directory Access Protocol

Active Directory clients must communicate with computers running Active Directory during logon to the network and when searching for shared resources. Access to domain controllers and global catalogs is performed using the Lightweight Directory Access Protocol (LDAP).

LDAP is a communication protocol designed for use on TCP/IP networks. LDAP defines how a directory client can access a directory server and how the client can perform directory operations and share directory data. LDAP standards are established by working groups of the Internet Engineering Task Force (IETF). Active Directory implements the LDAP attribute draft specifications and the IETF standards for LDAP versions 2 and 3.

As its name implies, LDAP is designed as an efficient method for accessing directory services without the complexity of other directory service protocols. Because LDAP defines what operations can be performed to query and modify information in a directory and how information in a directory can be securely accessed, you can use LDAP to find or enumerate directory objects and to query or administer Active Directory.

2.15.1 LDAP and Interoperability

LDAP is an open Internet standard. By using LDAP, Active Directory enables interoperability with other vendor directory services. Active Directory support for LDAP includes an LDAP provider object as part of Active Directory Service Interfaces (ADSI). ADSI supports the C-binding application programming interfaces for LDAP specified by Internet standard RFC 1823. Other directory service applications can be easily modified to access information in Active Directory by using ADSI and LDAP.

2.16 Single Master Operations

Active Directory supports multimaster replication of the directory data store between all domain controllers in the domain. Some changes are impractical to perform in multimaster fashion, however, so only one domain controller, called the *operations master*, accepts requests for such changes.

Since the operations master roles can be moved to other domain controllers within the domain or forest, these roles are sometimes referred to as flexible single master operations. In any Active Directory forest, there are five operations master roles that are assigned to one or more domain controllers. Some roles must appear in every forest. Other roles must appear in every domain in the forest.

2.16.1 Forest Wide Operations Master Roles

Every Active Directory forest must have the following roles:

1. Schema Master.
2. Domain Naming Master.

These roles must be unique in the forest. This means that throughout the entire forest there can be only one schema master and one domain naming master.

2.16.1.1 Schema Master

The schema master domain controller controls all updates and modifications to the schema. To update the schema of a forest, you must have access to the schema master. At any time, there can be only one schema master in the entire forest.

2.16.1.2 Domain Naming Master

The domain controller holding the domain naming master role controls the addition or removal of domains in the forest. There can be only one domain naming master in the entire forest at any time.

2.16.1.3 Domain Operations Master Roles

Every domain in the forest must have the following roles:

1. Relative ID Master.
2. Primary domain controller (PDC) emulator.
3. Infrastructure master.

These roles must be unique in each domain. This means that each domain in the forest can have only one relative ID master, PDC emulator, and infrastructure master.

2.16.2 Relative ID Master

The relative ID master allocates sequences of relative IDs to each of the various domain controllers in its domain. At any time, there can be only one domain controller acting as the relative ID master in each domain in the forest.

Whenever a domain controller creates a user, group, or computer object, it assigns the object a unique security ID. The security ID consists of a domain security ID (that is the same for all security IDs created in the domain), and a relative ID that is unique for each security ID created in the domain.

To move an object between domains (using Movetree.exe), you must initiate the move on the domain controller acting as the relative ID master of the domain that currently contains the object.

2.16.3 PDC Emulator

If the domain contains computers operating without Windows 2000 client software or if it contains Windows NT backup domain controllers (BDCs), the PDC emulator acts as a Windows NT primary domain controller. It processes password changes from clients and replicates updates to the BDCs. At any time, there can be only one domain controller acting as the PDC emulator in each domain in the forest.

In a Windows 2000 domain operating in native-mode, the PDC emulator receives preferential replication of password changes performed by other domain controllers in the domain. If a password was recently changed, that change takes time to replicate to every domain controller in the domain. If a logon authentication fails at another domain controller due to a bad password, that domain controller will forward the authentication request to the PDC emulator before rejecting the log on attempt.

2.16.4 Infrastructure Master

The infrastructure master is responsible for updating the group-to-user references whenever the members of groups are renamed or changed. At any time, there can be only one domain controller acting as the infrastructure master in each domain.

When you rename or move a member of a group (and that member resides in a different domain from the group), the group may temporarily appear not to contain that member. The infrastructure master of the group's domain is responsible for updating the group so it knows the new name or location of the member. The infrastructure master distributes the update via multimaster replication. There is no compromise to security during the time between the member rename and the group update. Only an administrator looking at that particular group membership would notice the temporary inconsistency.

2.17 Administering Active Directory

In an organization with more than one domain, it is often necessary to administer a domain other than the one to which you are currently logged on. For example, when creating trusts, the trust must be created in both the trusting and the trusted domain.

1. Cooperation with those who have administrative permissions in another domain.
2. Logging on with a user account with the necessary permissions.
3. Using the **Run as** command to run an administrative tool targeted on the particular domain (recommended).

A secure method of controlling administrative access to a particular domain is to very tightly control the number of accounts with administrative permissions for that domain and the number of people aware of those accounts. Only the people who know the account name and password can make administrative changes to the domain.

For example, if an administrator in another domain wanted to establish a shortcut trust with a tightly controlled domain, the only way that administrator could establish the trust relationship would be by communicating with the administrator of the tightly controlled domain, agreeing on a common password for the trust, and having the administrator of the tightly controlled domain create the trust in that domain.

A more convenient method of administering more than one domain is to log on with a user account that has administrative permissions in both domains. For example, user accounts that are members of the Enterprise Admins security group have permission to administer every domain in the forest. Using accounts with sweeping privileges is not recommended.

2.17.1 Delegating Administration

You can delegate administrative control to any level of a domain tree by creating organizational units within a domain and delegating administrative control for specific organizational units to particular users or groups. To decide what organizational units you want to create, and which organizational units should contain accounts or shared resources, consider the structure of your organization.

For example, you may want to create an organizational unit that enables you to grant to a user the administrative control for all user and computer accounts in all branches of a single organizational department, such as a Human Relations department. You may instead want to grant to a user administrative control only to some resources within a department, for example, computer accounts.

Another possible delegation of administrative control would be to grant to a user the administrative control for the Human Relations organizational unit, but not to any organizational units contained within the Human Relations organizational unit.

By delegating administrative responsibilities, you can eliminate the need for multiple administrative accounts that have broad authority (such as, over an entire domain). Although you likely will still use the predefined Domain Admins group for administration of the entire domain, you can limit the accounts that are members of the Domain Admins group to highly trusted administrative users. Windows 2000 defines many very specific permissions and user rights that can be used for the purposes of delegating or restricting administrative control.

Using a combination of organizational units, groups and permissions, you can define the most appropriate administrative scope for a particular person: an entire domain, all organizational units within a domain, or even a single organizational unit.

Administrative control can be granted to a user or group by using the Delegation of Control wizard. The Delegation of Control wizard allows you to select the user or group to which you want to delegate control, the organizational units and objects you want to grant those users the right to control and the permissions to access and modify objects. For example, a user can be given the right to modify the Owner of Accounts property, without being granted the right to delete accounts in that organizational unit.

2.17.1.1 Customizing MMC Consoles for Specific Groups

You can use MMC console options to create a limited-use version of a snap-in such as Active Directory Users and Computers. This allows administrators to control the options available to groups to whom you have delegated administrative responsibilities by restricting access to operations and areas within that customized console.

For example, suppose you delegate the Manage Printers right to the Print Managers group in the Manufacturing organizational unit. To simplify administration, you can create a custom console for use by members of the Print Managers group containing only the Manufacturing organizational unit and restrict the scope of the console using MMC console modes.

This type of delegation is also enhanced by the Group Policy settings available for MMC. These settings enable the administrator to establish which MMC snap-ins can be run by the affected user. The settings can be inclusive, allowing a set of snap-ins to run, or exclusive, restricting the set of snap-ins to run.

Using Group Policy you can distribute a customized console to specific groups in one of two modes: publishing or assigning. Publishing a customized console advertises the console to the members of a group specified in the Group Policy setting by adding the console to the list of available programs in Add/Remove Programs.

The next time the members of the group open Add/Remove Programs they have the option to install the new console. Assigning (as opposed to publishing) a console forces the console to be automatically installed for all specified accounts.

To publish or assign a console, create or modify a Group Policy object and apply it to the appropriate group of users. Then use the Software Installation extension of the Group Policy snap-in to either publish or assign the console.

1. The console must be packaged before using the Software Installation snap-in. You can use a tool such as Windows Installer to package the customized console. Once this has been accomplished you can configure the Software Installation snap-in to publish or assign the newly created package. For more information on how to package an application see the Windows 2000 Server SDK and Resource Kit.
2. If the customized console you are packaging uses a snap-in that is not installed on the destination workstation or server for the published or assigned user, you will need to include the snap-in file and the registration of the file in the package. You can either create a separate package that contains the snap-in or add the snap-in during the creation of the customized console package so that it will be properly installed on the computer every time a user installs the console package.

2.17.2 Operations on Master Failures

Some of the operations master roles are crucial to the operation of your network. Others can be unavailable for quite some time before their absence becomes a problem. Generally, you will notice that a single master operations role holder is unavailable when you try to perform some function controlled by the particular operations master.

If an operations master is not available due to computer failure or network problems, you can seize the operations master role. This is also referred to as forcing the transfer of the operations master role. Before forcing the transfer, first determine the cause and expected duration of the computer or network failure. If the cause is a networking problem or a server failure that will be resolved soon, wait for the role holder to become available again.

If the domain controller that currently holds the role has failed, you must determine if it can be recovered and brought back online. In general, seizing an operations master role is a drastic step that should be considered only if the current operations master will never be available again. The decision depends upon the role and how long the particular role holder will be unavailable. The impact of various role holder failures is discussed in the following topics.

2.17.2.1 Schema Master Failure

Temporary loss of the schema operations master is not visible to network users. It will not be visible to network administrators either, unless they are trying to modify the schema or install an application that modifies the schema during installation.

If the schema master will be unavailable for an unacceptable length of time, you can seize the role to the standby operations master. However, seizing this role is a drastic step that you should take only when the failure of the schema master is permanent.

2.17.2.2 Domain Naming Master Failure

Temporary loss of the domain naming master is not visible to network users. It will not be visible to network administrators either, unless they are trying to add a domain to the forest or remove a domain from the forest.

If the domain naming master will be unavailable for an unacceptable length of time, you can seize the role to the standby operations master. However, seizing this role is a drastic step that you should take only when the failure of the domain naming master is permanent.

2.17.2.3 Relative ID Master Failure

Temporary loss of the relative identifier operations master is not visible to network users. It will not be visible to network administrators either, unless they are creating objects and the domain in which they are creating the objects runs out of relative identifiers. If the relative identifier master will be unavailable for an unacceptable length of time, you can seize the role to the operations master.

However, seizing this role is a drastic step that you should take only when the failure of the relative identifier master is permanent.

2.17.2.4 PDC Emulator Failure

The loss of the primary domain controller (PDC) emulator affects network users. Therefore, when the PDC emulator is not available, you may need to immediately seize the role.

If the current PDC emulator master will be unavailable for an unacceptable length of time and its domain has clients without Windows 2000 client software, or if it contains Windows NT backup domain controllers, seize the PDC emulator master role to the standby operations master. When the original PDC emulator master is returned to service, you can return the role to the original domain controller.

2.17.2.5 Infrastructure Master Failure

Temporary loss of the infrastructure master is not visible to network users. It will not be visible to network administrators either, unless they have recently moved or renamed a large number of accounts.

If the infrastructure master will be unavailable for an unacceptable length of time, you can seize the role to a domain controller that is not a global catalog but is well connected to a global catalog (from any domain), ideally in the same site as the current global catalog. When the original infrastructure master is returned to service, you can transfer the role back to the original domain controller.

2.17.3 Service Duplication

The following topics describe some types of service information that may be useful to publish to the directory. The qualities that make a service appropriate for publishing may be better understood by understanding how Active Directory uses services.

2.17.3.1 Service Categories

Binding and configuration information are the two types of information frequently published using Active Directory Binding information allows clients to connect to services that do not have well known bindings and that conform to a service-centric model.

By publishing the bindings for these kinds of services, Windows 2000 can automatically establish connections with services. Machine-centric services are typically handled on a service-by-service basis and should not be published to the directory.

Configuration information can be common across client applications. Publishing this sort of information allows you to distribute current configuration information for these applications to all clients in the domain. The configuration information is accessed by client applications as needed. This eases application configuration for users and gives you more control over application behaviors.

2.17.3.2 Service Information Characteristics

Service information that you publish to the directory is most effective if it has the following characteristic:

1. **Useful to Many Clients:** Information that is useful to a small set of clients or that is useful only in certain areas of the network should not be published. If not widely used, this information wastes network resources, since it is published to every domain controller in the domain.
2. **Relatively Stable and Unchanging:** Although there may be exceptions to this rule, it generally makes sense to publish only service information that changes less frequently than two replication intervals. For intra-site replication, the maximum replication period is fifteen minutes, and for inter-site replication, the maximum replication period is configured based on the replication interval of the site link used for the replication. Object properties that change more frequently create excessive demands on network resources. Property values may be out-of-date until updates are published, which can take as long as the maximum replication period. Consequently, having properties out-of-date

for that period of time must not create unacceptable conditions. For example, some network services select a valid TCP port for use each time they are started. After selecting the port, the service updates Active Directory with this information, which is stored as the service connection point. Clients access the service connection point when they want to use the service, but if the new service connection point has not been replicated when the client requests it; the client will receive an outdated port, rendering the service temporarily inaccessible.

3. **Well-Defined, Reasonable Properties:** Information that is of a consistent form is easier for services to use. The information should be relatively small in size.

2.17.4 Managing Security

Windows 2000 introduces several new security features to help you implement the level of security that your organization needs. In its simplest form, security ensures that the people logging on to your network are who they say they are.

When you create trust relationships between Windows domains or Kerberos V5 realms, Windows security can limit access to sensitive data or specific resources to only those people to whom you want to grant access, both within and outside your organization. Windows 2000 security can also ensure that the data you store on disk, or send over private or public networks is protected from unauthorized access. You can use Encrypting File System to protect data stored on disk. IP security and PPTP encryption can protect data on your network as well as data transmitted over the Internet.

2.17.5 Programming Interfaces

Active Directory Service Interfaces (ADSI) provides a simple, powerful, object-oriented interface to Active Directory. ADSI makes it easy for programmers and administrators to create directory programs by using high-level tools such as Microsoft Visual Basic, Java, C, or Visual C++ without having to worry about the underlying differences between the different namespaces.

ADSI enables you to build or buy programs that give you a single point of access to multiple directories in your network environment, whether those directories are based on LDAP or another protocol. ADSI is fully scriptable for ease of use by administrators.

Active Directory also provides support for Messaging Application Programming Interface (MAPI), so legacy MAPI programs will continue to work with Active Directory. In addition, Active Directory supports the LDAP C API, defined in RFC 1823, as a lower-level interface for C programmers.

2.17.6 Active Directory Administrative Tools

The Active Directory administrative tools that are included with Windows 2000 Server simplify directory service administration. You can use the standard tools or, using Microsoft Management Console (MMC), create custom tools that focus on single management tasks. You can combine several tools into one console. You can also assign custom tools to individual administrators with specific administrative responsibilities.

The Active Directory administrative tools can only be used from a computer with access to a Windows 2000 domain. The following Active Directory administrative tools are available on the Windows 2000 Server Administrative Tools menu of all Windows 2000 domain controllers:

1. Active Directory Users and Computers.
2. Active Directory Domains and Trusts.
3. Active Directory Sites and Services.

You can also administer Active Directory remotely, from a computer that is not a domain controller. To use the Active Directory administrative tools remotely, from a computer that is not a domain controller, such as one running Windows 2000 Professional, you must install the Windows 2000 Administrative Tools.

The Active Directory Schema snap-in is another Active Directory administrative tool. It is not available on the Windows 2000 Server Administrative Tools menu. You must install the Windows 2000 Administration Tools from the Windows 2000 Server compact disc and add it to an MMC console. For advanced administrators and network support specialists, there are many command line tools that can be used to configure, manage, and troubleshoot Active Directory. You can also create scripts that use Active Directory Service Interfaces (ADSI). Several sample scripts are supplied on the Windows 2000 Server compact disc.

Chapter III: Network Security

3.1 Introduction

The primary features of the Microsoft Windows Server family security model are user authentication and access control. The Active Directory service ensures that administrators can manage these features easily and efficiently.

3.2 Authentication

Interactive logon confirms the user's identification to the user's local computer or Active Directory account.

Network authentication confirms the user's identification to any network service that the user is attempting to access. To provide this type of authentication, the security system includes these authentication mechanisms: Kerberos V5, public key certificates, Secure Sockets Layer/Transport Layer Security (SSL/TLS), Digest, and NTLM (for compatibility with Windows NT 4.0 systems).

Authentication in the Windows Server family also includes two-factor authentication, such as smart cards.

3.3 Object-Based Access Control

Along with user authentication, administrators are allowed to control access to resources or objects on the network. To do this, administrators assign security descriptors to objects that are stored in Active Directory. A security descriptor lists the users and groups that are granted access to an object and the specific permissions assigned to those users and groups. A security descriptor also specifies the various access events to be audited for an object. Examples of objects include files, printers, and services. By managing properties on objects, administrators can set permissions, assign ownership, and monitor user access.

Not only can administrators control access to a specific object, they can also control access to a specific attribute of that object. For example, through proper configuration of an object's security descriptor, a user could be allowed to access a subset of information, such as employees' names and phone numbers but not their home addresses.

3.4 Security Policy

You can control security on your local computer or on multiple computers by controlling password policies, account lockout policies, Kerberos policies, auditing policies, user rights, and other policies. To create a system wide policy, you can use security templates, apply templates using Security Configuration and Analysis or edit policies on the local computer, organizational unit, or domain.

3.5 Auditing

Monitoring the creation or modification of objects gives you a way to track potential security problems, helps to ensure user accountability, and provides evidence in the event of a security breach.

3.6 Security Measures for Active Directory

Active Directory provides protected storage of user account and group information by using access control on objects and user credentials. Because Active Directory stores not only user credentials but also access control information, users who log on to the network obtain both authentication and authorization to access system resources. For example, when a user logs on to the network, the security system authenticates the user with information stored in Active Directory. Then, when the user attempts to access a service on the network, the system checks the properties defined in the discretionary access control list (DACL) for that service.

Because Active Directory allows administrators to create group accounts, administrators can manage system security more efficiently. For example, by adjusting a file's properties, an administrator can permit all users in a group to read that file. In this way, access to objects in Active Directory is based on group membership.

3.6.1 Trusts

The Windows Server family supports domain trusts and forest trusts. Domain trust allows a user to authenticate to resources in another domain. In a Windows Server forest, administrators can create a forest to extend two-way transitivity beyond the scope of a single forest to a second Windows Server 2003 forest.

3.6.2 Access Control

In order to secure a computer and its resources, you must take into consideration what rights users will have. You can secure a computer or multiple computers by granting users or groups specific user rights. You can secure an object, such as a file or folder, through assigning permissions to allow users or groups to perform specific actions on that object.

Access control is the process of authorizing users, groups, and computers to access objects on the network. Key concepts that make up access control are permissions, user rights, and object auditing.

3.6.3 Permissions

Permissions define the type of access granted to a user or group for an object or object property. For example, the Finance group can be granted Read and Write permissions for a file named Payroll.dat.

Permissions are applied to any secured objects such as files, Active Directory objects, or registry objects. Permissions can be granted to any user, group, or computer. It is a good practice to assign to groups.

The permissions attached to an object depend on the type of object. For example, the permissions that can be attached to a file are different from those that can be attached to a registry key. Some permissions, however, are common to most types of objects. These common permissions are:

- Read permissions
- Modify permissions
- Change owner
- Delete

When you set up permissions, you specify the level of access for groups and users. For example, you can let one user read the contents of a file, let another user make changes to the file, and prevent all other users from accessing the file. You can set similar permissions on printers so that certain users can configure the printer and other users can only print from it.

3.6.4 Inheritance of Permissions

Inheritance allows administrators to easily assign and manage permissions. This feature automatically causes objects within a container to inherit all the inheritable permissions of that container. For example, the files within a folder, when created, inherit the permissions of the folder. Only permissions marked to be inherited will be inherited.

3.7 Authorization Manager

Authorization Manager provides a flexible framework for integrating role–based access control into applications. It enables administrators who use those applications to provide access through assigned user roles that relate to job functions. Authorization Manager applications store authorization policy in the form of authorization stores that are stored in Active Directory or XML files and apply

3.8 Security Configuration Manager

The Security Configuration Manager tool set allows you to create, apply and edit the security for your local computer, organizational unit, or domain.

3.8.1 Auditing Security Events

You can set up audit policy so that user or system activity in specified event categories is recorded. You can monitor security-related activity, such as who accesses an object, if a user logs on to or logs off from a computer, or if changes are made to an auditing policy setting.

3.9 Encrypting File System

With Encrypting File System (EFS), you can encrypt files and directories that are stored on a disk. Encrypting File System (EFS) provides the core file encryption technology used to store encrypted files on NTFS file system volumes. Once you encrypt a file or folder, you work with the encrypted file or folder just as you do with any other files and folders.

Encryption is transparent to the user that encrypted the file. This means that you do not have to manually decrypt the encrypted file before you can use it. You can open and change the file as you normally do.

Using EFS is similar to using permissions on files and folders. Both methods can be used to restrict access to data. However, an intruder who gains unauthorized physical access to your encrypted files or folders will be prevented from reading them. If the intruder tries to open or copy your encrypted file or folder he receives an access denied message. Permissions on files and folders does not protect against unauthorized physical attacks.

You encrypt or decrypt a folder or file by setting the encryption property for folders and files just as you set any other attribute such as read-only, compressed, or hidden. If you encrypt a folder, all files and subfolders created in the encrypted folder are automatically encrypted. It is recommended that you encrypt at the folder level.

When you work with encrypted files and folders, keep in mind the following information:
Only files and folders on NTFS volumes can be encrypted. Because WebDAV works with NTFS, NTFS is required when encrypting files over WebDAV (Web distributed authoring and versioning).

Files or folders that are compressed cannot also be encrypted. If the user marks a file or folder for encryption, that file or folder will be uncompressed. Encrypted files can become decrypted if you copy or move the file to a volume that is not an NTFS volume.

Moving unencrypted files into an encrypted folder will automatically encrypt those files in the new folder. However, the reverse operation will not automatically decrypt files. Files must be explicitly decrypted. Files marked with the System attribute cannot be encrypted, nor can files in the system root directory structure.

Encrypting a folder or file does not protect against deletion or listing files or directories. Anyone with the appropriate permissions can delete or list encrypted folders or files. For this reason, using EFS in combination with NTFS permissions is recommended.

You can encrypt or decrypt files and folders located on a remote computer that has been enabled for remote encryption. However, if you open the encrypted file over the network, the data that is transmitted over the network by this process is not encrypted. Other protocols, such as SSL/TLS (Secure Socket Layer/Transport Layer Security) or Internet Protocol security (IPSec) must be used to encrypt data over the wire. WebDAV, however, is able to encrypt the file locally and transmit it in encrypted form.

3.9.1 Encrypting and Decripting Data

With Encrypting File System (EFS) you can store data securely. EFS does this by encrypting data in selected NTFS file system files and folders.

Because EFS is integrated with the file system, it is easy to manage, difficult to attack, and transparent to the user. This is particularly useful for securing data on computers that may be vulnerable to theft, such as mobile computers.

Files and folders cannot be encrypted or decrypted on FAT volumes. Also, EFS is designed to store data securely on local computers. As such, it does not support the secure transmission of files over a network. Other technologies, such as Internet Protocol Security (IPSec), can be used in conjunction with EFS to provide a larger solution.

3.9.2 Using Encryption Keys

Once a user has specified that a file be encrypted, the actual process of data encryption and decryption is completely transparent to the user. The user does not need to understand this process. However, the following explanation of how data encryption and decryption works might be useful for administrators.

3.9.3 Encryption of Files

Each file has a unique file encryption key, which is later used to decrypt the file's data. The file encryption key is itself encrypted, it is protected by the user's public key corresponding to the user's EFS certificate.

The file encryption key is also protected by the public key of each additional EFS user that has been authorized to decrypt the file and each recovery agent.

The EFS certificate and private key used can be issued by a number of sources, including automatically-generated certificates, certificates created by Microsoft certification authorities (CAs), or third-party CAs.

3.9.4 Decryption of Files

To decrypt a file, the file encryption key must first be decrypted. The file encryption key is decrypted when the user has a private key that matches the public key.

The original user may not be the only person that can decrypt the file encryption key. Other designated users or recovery agents can also decrypt the file encryption key, by using their own private key.

3.9.5 Storing Encrypted Files on A Remote Server

If users in your Windows XP or Windows Server family computing environment want to store encrypted files on remote servers, it is useful to know the following:

Windows XP and the Windows Server family support the storage of encrypted files on remote servers. Users can use EFS remotely only when both computers are members of the same Windows Server family forest.

Encrypted data is not encrypted when in transit over the network, but only when stored on disk. The exceptions to this are when your system includes Internet Protocol security (IPSec) or Web Distributed Authoring and Versioning (WebDAV). IPSec encrypts data while it is transported over a Transmission Control Protocol/Internet Protocol (TCP/IP) network. If the file is encrypted before being copied or moved to a WebDAV folder on a server, it will remain encrypted during the transmission and while it is stored on the server.

Encrypted files are not accessible from Macintosh clients. Storing EFS certificates and private keys on smartcards is not currently supported. Strong private key protection for EFS private keys is not currently supported.

Before users can encrypt files that reside on a remote server, an administrator must designate the remote server as trusted for delegation. This allows all users with files on that server to encrypt those files.

3.10 Recording Data

Data recovery is important when you need to be able to recover data encrypted by an employee after the employee leaves, or when the user loses the private key. Data recovery is available through the Encrypting File System (EFS) as a part of the overall security policy for the system. For example, if you should ever lose your file encryption certificate and associated private key through disk failure, arson, or any other reason, the person who is the designated recovery agent can recover the data. In a business environment, an organization can recover data encrypted by an employee after the employee leaves.

3.11 Recovery Policy

EFS uses recovery policies to provide built-in data recovery. A recovery policy is a type of public key policy that provides for one or more user accounts to be designated as recovery agents.

A recovery policy is configured locally for stand-alone computers. For computers that are part of a network, a recovery policy is configured at the domain, organizational unit, or individual computer level, and applies to all Windows XP and Windows Server family-based computers that the policy applies to. A certification authority (CA) issues recovery certificates, and you use Certificates in Microsoft Management Console (MMC) to manage them.

In a domain, the Windows Server family implements a default recovery policy for the domain when the first domain controller is set up. The self-signed certificate is issued to the domain administrator. That certificate designates the domain administrator as the recovery agent. To change the default recovery policy for a domain, log on to the first domain controller as an administrator. Additional recovery agents can be added to this policy and the original recovery agent can be removed at any time.

Because the Windows XP and Windows Server family security subsystems handle enforcing, replicating, and caching of the recovery policy, users can implement file encryption on a system that is temporarily offline, such as a portable computer. This process is similar to logging on to their domain account using cached credentials.

3.12 Recovery Agents

A recovery agent is an individual authorized to decrypt data that was encrypted by another user. Recovery agents do not need any other permissions to function in this role. Recovery agents are useful, for example, when employees leave the company and their remaining data needs to be decrypted. Before you can add a recovery agent for a domain, you must ensure that each recovery agent has been issued an X.509v3 certificate.

Each recovery agent has a special certificate and associated private key that allows data recovery wherever the recovery policy applies. If you are the recovery agent, you should be sure to use the Export command in Certificates in MMC to back up the recovery certificate and the associated private key to a secure location. After backing them up, you should use Certificates in MMC to delete the recovery certificate. Then, when you need to perform a recovery operation for a user, you should first restore the recovery certificate and associated private key using the Import command from Certificates in MMC. After recovering the data, you should again delete the recovery certificate. You do not have to repeat the export process.

To add recovery agents for a domain, you add their certificates to the existing recovery policy. For steps on how to add recovery agents to a domain, see To add a recovery agent for a domain.

3.13 Managing Certificates

Encrypting File System (EFS) uses public key cryptography to encrypt the contents of files. The keys that are used are obtained from the certificate of the user and any additional users and designated recovery agents configured. Because the certificates may also contain private key information, they must be managed correctly.

Certificates that are used by EFS can be obtained from a certification authority (CA) or created automatically by the computer. When obtaining an EFS certificate from a CA, the cryptographic service provider (CSP) and the appropriate object identifier (also known as an OID) must be referenced by the certificate. EFS can use either a base or enhanced CSP. If these two attributes are not set correctly in the certificate, EFS is unable to use it. The certificate and private key of all designated recovery agents should be exported to removable disk and stored securely until needed.

When exporting the certificate and private key, ensure that, in the Intended Purposes column, the selected certificate includes Encrypting File System and that you have the associated private key.

Certificates and private keys can be used on multiple computers. If your Windows XP network is configured to use roaming profiles, the certificates will be available on any computer you logon to. Otherwise, the certificates and private keys must be exported and imported manually. To manually export a certificate and private key, use the above procedure.

3.13.1 Verificaion of Certificate Validity

Certificates are not expected to be valid indefinitely. Over time, an attacker can determine the corresponding private key and render the data encrypted with that key vulnerable. Because of this, certificates have a validity period that defines the length of time during which they can be considered valid. Once the validity period expires, a new certificate must be obtained to encrypt new data. However, the existing certificate and private key are normally retained so older data can still be decrypted.

Certificates can also be revoked by the issuing CA, even when they are still within their validity period. There are a number of reasons why a certificate could become untrustworthy, including compromise of the certificate subject's private key or discovery that a certificate was obtained fraudulently. When a certificate is revoked, it is placed on a certificate revocation list (CRL) maintained by the CA.

When a file is encrypted, EFS checks the validity period on the certificate of the user as well as the recovery agent. When a new user is added to an existing file, EFS checks for both the revocation of the certificate being added and the chaining of the certificate to a trusted root CA. If the certificate is found to be invalid (either because of expiration, revocation or inability to chain) the certificate is not used and the user is typically notified.

3.13.2 Public Key Infrastructure

A public key infrastructure, often shortened to PKI, is a system of digital certificates, certification authorities (CAs) and other registration authorities (RAs) that verify and authenticate the validity of each party involved in an electronic transaction through the use of public key cryptography.

Internet Protocol security (IPSec) provides the following new features for enhanced security, scalability, and availability, and ease of deployment and administration.

3.14 IP Security Monitor

In Windows 2000, IP Security Monitor was implemented as an executable program (IPSecmon.exe). In Windows XP and the Windows Server family, IP Security Monitor is implemented as a Microsoft Management Console (MMC) console and includes enhancements that allow you to:

- Monitor IPSec information for your local computer and for remote computers.
- View details about active IPSec policies, including the name, description, date last modified, store, path, organizational unit, and Group Policy object name.
- View main mode and quick mode generic filters and specific filters.
- View main mode and quick mode statistics. For information about the statistics displayed in IP Security Monitor, see Viewing main mode and quick mode statistics in IP Security Monitor.
- View main mode and quick mode security associations.

Customize refresh rates, and use DNS name resolution for filter and security association output. Search for specific main mode or quick mode filters that match any source or destination IP address, a source or destination IP address on your local computer, or a specific source or destination IP address.

3.15 Stronger Cryptopgraphic Master Key (Diffie-Hellman)

For enhanced security, IPSec now supports the use of a 2048-bit Diffie-Hellman key exchange. With a stronger Diffie-Hellman group, the secret key that is derived from the Diffie-Hellman exchange has greater strength. Strong Diffie-Hellman groups combined with longer key lengths increase the computational difficulty of determining a secret key.

3.16 Startup Security

For enhanced security, IPSec now provides stateful filtering of network traffic during computer startup. With stateful filtering, only the following traffic is permitted during computer startup: the outbound traffic that the computer initiates during startup, the inbound traffic that is sent in response to the outbound traffic, and DHCP traffic. As an alternative to stateful filtering, you can specify that all inbound and outbound traffic be blocked until an IPSec policy is applied. If you use stateful filtering, or if you specify that traffic be blocked during computer startup, you can also specify the traffic types that you want to exempt from IPSec filtering during computer startup.

3.17 Persistent Policy for Enhanced Security

You can now create and assign a persistent IPSec policy to secure a computer if a local IPSec policy or an Active Directory-based IPSec policy cannot be applied. When you create and assign a persistent policy, it is applied before the local policy or the Active Directory-based policy is applied, and it remains in effect regardless of whether the local policy or the Active Directory-based policy is applied (for example, an IPSec policy will not be applied if it is corrupted).

3.17.1 IPSEC Certificate to Account Mapping for Network Access Control

With the Windows Server family, if you use either Kerberos V5 or certificate authentication, you can set restrictions on which computers are allowed to connect. This functionality allows you to use IPSec to allow or deny any of the following access to a server running Windows Server :

- Computers that are members of a specific domain.
- Computers that have a certificate from a specific issuing certification authority.
- A specific group of computers.
- A specific computer.

When you enable certificate to account mapping in IPSec, the IKE protocol associates (maps) a computer certificate to a computer account in an Active Directory domain or forest, and then retrieves an access token, which includes the list of the user rights that are assigned to the computer.

You can restrict access by configuring Group Policy security settings and assigning either the Access this computer from the network user right or the Deny access to this computer from the network user right to individual or multiple computers, as needed.

Ability to exclude the name of the certification authority (CA) from certificate requests. For enhanced security, when you use certificate authentication to establish trust between IPSec peers, you can now exclude the name of the CA from the certificate request. When you exclude the name of the CA from the certificate request, you prevent the potential disclosure of sensitive information about the trust relationships of a computer, such as name of the company that owns the computer and the domain membership of the computer (if an internal public key infrastructure is being used), to an attacker.

3.17.2 IPSEC Policy Filters

You can now use the IP Security Policy Management console to configure the source or the destination address fields that the local IPSec policy will interpret as the addresses for the DHCP server, the DNS servers, the WINS servers, and the default gateway. As a result, IPSec policies can now automatically accommodate changes in the IP configuration of the server, by using either DHCP or static IP configurations.

3.17.3 IPSEC Functionality Over Network Address Translation (NAT)

IPSec Encapsulating Security Payload (ESP) packets can now pass through NATs that allow User Datagram Protocol (UDP) traffic. The IKE protocol automatically detects the presence of a NAT and uses UDP-ESP encapsulation to allow IPSec traffic to pass through the NAT. This functionality is an implementation of the Internet Engineering Task Force (IETF) IP Security Working Group standard for IPSec.

NATs are widely used for Internet Connection Sharing (ICS) and in locations that provide public Internet access (such as hotels and airports) and that are likely to be used by telecommuters. In addition, some Internet service providers (ISPs) use a centralized NAT to connect their clients to the Internet.

IPSec functionality over NAT enables IPSec-secured connections to be established in the following common deployment scenarios:

Layer Two Tunneling Protocol (L2TP)/IPSec virtual private network (VPN) clients that are behind NATs can establish IPSec-secured connections over the Internet to their corporate network, using IPSec ESP transport mode.

Servers running Routing and Remote Access can establish gateway-to-gateway IPSec tunnels when one of the servers running Routing and Remote Access is behind a NAT.

Clients and servers can send IPSec-secured TCP and UDP packets to other clients or servers using IPSec ESP transport mode, when one or both of the computers are behind a NAT. For example, a program running on a server on a perimeter network can be IPSec-secured when it is used to make connections to the corporate network.

3.17.4 Improved IPSEC Integration with Network Load Balancing

Improved IPSec integration with Network Load Balancing allows a Network Load Balancing group of servers to provide highly available IPSec-based VPN services. Network Load Balancing can accurately track IPSec-secured sessions, and the IPSec IKE protocol can detect when an IPSec-secured session is being established with a cluster server and quickly recover from a failover. Additionally, Network Load Balancing can now maintain IPSec-secured connections to the correct Network Load Balancing host, even when the number of hosts in the cluster (and the algorithm used to map clients to hosts) changes. Because the IKE protocol automatically detects the Network Load Balancing service, no additional configuration is required to use this feature.

3.17.5 IPSEC Support for RSOP

To enhance IPSec deployment and troubleshooting, IPSec now provides an extension to the Resultant Set of Policy (RSoP) console. RSoP is an addition to Group Policy that you can use to view existing IPSec policy assignments for a computer or for members of a Group Policy container. To view IPSec policy assignments for a computer, run an RSoP logging mode query.

To view IPSec policy assignments for members of a Group Policy container, run an RSoP planning mode query. For information about how to view IPSec policy assignments, see Use Resultant Set of Policy (RSoP) to view IPSec policy assignments.

After you run an RSoP logging mode query or an RSoP planning mode query, you can view detailed settings (the filter rules, filter actions, authentication methods, tunnel endpoints, and connection types that were specified when the IPSec policy was created) for the IPSec policy that is being applied.

Summary and Conclusion

Windows 2000 Advanced Server includes all the new features of Windows 2000 Server, and in addition offers enhanced memory support, support for additional processors, and clustering. Enhanced memory and processor support means your server applications can faster, providing better response for users on the network. Windows 2000 Professional offers increased compatibility with different types of networks and with a wide array of legacy hardware and software.

Active Directory is the directory service for Windows 2000 Server. It stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory directory service uses a structured data store as the basis for a logical, hierarchical organization of directory information.

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

Groups are Active Directory or local computer objects that can contain users, contacts, computers, and other groups. Distribution groups can only be used as e-mail distribution lists. They cannot be used to filter Group Policy settings. Distribution groups have no security function. As opposed to groups, organizational units are used to create collections of objects within a single domain, but do not confer membership. The administration of an organizational unit and the objects it contains can be delegated to an individual.

Many research centers and University Campuses require using advanced level of Network functionality, management, design and structure. There are many factors that need to be considered in organization of Network Management.

References

1. James Martin, "Computer *Networks and Distributed Processing*", Prentice-Hall, 1994, ISBN: 81-203-0529-9.
2. Andrew S. Tanenbaum, "*Computer Networks*", Second Edition, Prentice-Hall International, 1989, ISBN: 0-13-166836.
3. Dervis Z. Deniz, "ISDN *and Its Applications to LAN Interconnection*", Mc-Graw Hill, 1994, ISBN: 0-07-707883-7.