

Operating Systems with Windows NT2000

Course Designer and Acquisition Editor

Centre for Information Technology and Engineering

Manonmaniam Sundaranar University

Tirunelveli

CONTENTS

Lecture	Introduction to Networks	1
	Introduction to Networks	
	About PC Network	
	Concept of Networking	
	Benefits of Networks	
	Types of Networks	
	Classification of Networks	
Lecture 2	Introduction to NT	18
	Introduction to NT	
	Networking with Windows NT Server 4.0	
	Windows NT Server 4.0, the Internet and Intranets	
	What New in Windows NT Server 4.0	
	Coming Attraction in Windows NT 5.0	
Lecture 3	Understanding the Windows NT Operating System	27
	Windows NT Operating System Features	
	Windows NT System Architecture	
Lecture 4	NT Environmental Subsystems	38
	Client and Protected subsystems server	
Lecture 5	Choosing Network Protocol	46
	Understand the OSI Seven - Layer Model	
	Comparing Windows NT and OSI Network Layers	
	Networking with Windows NT's Protocols	
	Supporting a Variety of PC Clients	
Lecture 6	Network Topologies & Architecture	63
	Access Methods	
	Network Topologies	
	Network Architecture	

Lecture 7	Transmission Media	80
	Cable Media Wireless Media Comparisons of different Wireless Media	
Lecture 8	Network Adapter Card	97
	Working of a Network Adapter Card Network Adapter Card Compatibility Configuring Network Adapter Cards	
Lecture 9	Connectivity Devices and Transfer Mechanism	109
	Addressing Modems Repeaters Hubs Bridges Routing Gateways	
Lecture 10	File systems	123
	File System NTFS File System Understanding NTFS Permission Compressing NTFS Files and Folders	
Lecture 11	Computer Security	143
	Security & Windows NT C2 Security Requirements For C2 Security	
Lecture 12	Windows NT Server Installation	149
	System Requirements Compatibility Issues Types of Installation Starting the Basic Installation Repairing the Windows NT Server Operating System Installation	

Lecture 13	Setting up RAID	179
	Understanding RAID Levels	
	Creating Windows NT Server Stripe and Mirror sets	
	Recovering a Software RAID 1 or RAID 5 Set	
Lecture 14	Installing File Backup Systems	189
	Backup Types	
	Developing a Backup Strategy	
	Choosing Backup Hardware	
	Windows NT Server 4.0 Backup Application	
Lecture 15	Windows NT Registry	204
	Registry Basics	
	Configuration settings in Registry	
	Registry's Organization	
	Registry Editor	
	Important Hives & Keys	
	Inspecting Another Computer's Registry	
	Maintaining Registry Security	
Lecture 16	Using TCP/IP, WINS and DHCP	224
	Role of TCP/IP	
	Installing & Configuration TCP/IP	
	Implementing DHCP	
	Implementing WINS	
Lecture 17	Working with Domains	225
	Win NT- Domain Models	
	Domain Architecture & Security	
	Implementing Domains and Trusts between Domains	
Lecture 18	Managing User Account	265
	Working with User Manger for Domains	
	Managing User accounts & their Properties	
	Administering the Domain Account Policy	

Lecture 19	Managing Group Accounts	288
	Managing User Groups Using Group Management Wizard Managing User Rights Policy	
Lecture 20	Sharing and Securing Network Resources	300
	Sharing & Securing Folders and Files Replication Folders Sharing and Securing Network Printers	
Lecture 21	Monitoring the Network	321
	Performance Monitor Network Monitor Event Viewer and Log Files	
Lecture 22	Optimizing the Network Server Performance	3440
	Optimizing NT File and Print Server Optimizing NT as an Application Server	
Lecture 23	Troubleshooting	356
	Hardware Problems Boot Failure Relating Network Protocols and Troubleshooting Issues Network Problem Using Protocol Analyzers Windows NT Server 4.0's Primary Troubleshooting tools	
Lecture 23	Syllabus	374

Lecture 1

Introduction to Networks

Objectives

In this Lecture you will learn the following:

- Understanding the concept of network and its benefit
- Classification Of Networking
- Knowing the Types Of Network

Coverage Plan

Lecture 1

- 1.1 Snap Shot
- 1.2 About PC Network
- 1.3 Concept of Networking
- 1.4 Benefits Of Networking
- 1.5 Types Of Networking
- 1.6 Classification Of Networking
 - 1.6.1 Peer-to-peer-Networks
 - 1.6.1.2 Peer-to-peer Operating Systems
 - 1.6.2 Server-Based Network
- 1.7 Short Summary
- 1.8 Brain Storm

1.1 Snap Shot - What is networking?

A Computer network is a communication system, which connects two or more computers and peripheral devices, to allow sharing of resources, information, and services. Users in a network share a common pathway to communicate with each other.

Connecting a number of computers (server and workstations), using wired media (cables) or wireless media forms a computer network. Each computer has a network interface card (NIC) installed in it and runs a piece of software (called network software) to achieve connectivity.

Definition

“A group of computer that are connected together for the purpose of sharing various computer software and hardware resources like input-output devices, memory, storage devices like hard disk etc.”

1.2 About PC Networks

Before 1980s, mainframes and minicomputers were popular. These systems were, expensive, bulky, huge in size, and located in special rooms. The programs as well as data were centrally stored on the host (server) and processed centrally by the host. The clients had to work on “dumb’ terminals, which had no processing power. The terminals were connected to the host using a serial cable, supporting a data transfer rate of 9600 bits per second or lower. This type of processing is called **centralized processing**. It puts an extra burden on the host, as the host not only stores the data but also processes the data for every individual terminal. Though the host is a very high performance server, the centralized processing mechanism easily overloads the host even with the addition of few more users. This drastically reduces the overall network performance. Another disadvantage of mainframe systems was, they were proprietary to vendors, and hence the system expansion was difficult as well as expensive.

Centralized Processing

In Centralized processing the applications and data are centrally stored on the host and shared by the terminals attached to the host. The important aspect to be noticed here is that, the terminals don’t have any processing power and the data storage as well as processing is done on the host itself.

Pros

- As the data storage as well as processing is done on the hosts systems, it helps in reducing the cost of hardware for the client systems.

Cons

- The cost of hardware required for the host is very high as it needs very high capacity hard disk to store the shared applications and user data, and needs large amount of RAM as well.
- The host gets loaded with additional number of terminals, as the terminals do not have the processing power of their own.

Emergence of PC networks

The first half of 1980's saw the rise of personal computers. The Personal Computers (PCs) were restricted to a single user operating system (DOS) in its early days, and hence did not feature networking. Each PC was required to have its own resources (which could not be shared), and the transfer of data between the PCs was through floppy diskettes.

As a first PC networking exercise, the UNIX operating system was ported to PCs, this worked fine. Later many network operating systems were designed, which broke the mainframe and minicomputer users to desktops. Novell's NetWare operating system gave popularity to PC networks and was the most widely used network operating system until 1996/97. Today Microsoft's Windows NT is leading the show with the largest market share among all the available operating systems.

Except UNIX, the other entire range of network operating system used with PCs (E.g. Novell's NetWare, Microsoft's Windows NT, and IBM's OS/2), feature distributed processing. In distributed processing environment, the programs (may be the data) are centrally stored on the server, but the processing is done individually at each workstation. This means when a user working on his PC (called workstation) requests for the program (or data) the request is sent to the server (usually called a file server.) The server then gives out the requested program (or data) to that workstation. When the user processes this data, the workstation does all the processing. This is different from the centralized processing, in which the data is processed at the host itself (or server) and only the result is displayed on the terminal.

Distributed Processing

In distributed processing system, each workstation processes its own data, the server does not get loaded easily with the addition of users, as in the case with centralized

processing. However, unlike the cheaper dumb terminals used in centralized environments, the distributed environments require every machine to have the required hardware.

- As the data is processed by individual workstation, the server does not get easily loaded with addition of more workstation.
- The server need not have very high amount of memory, and hard disk capacity(if workstation have hard disks)
- Each workstation should have its own high performance processor, and memory that increases the hardware cost.

1.3 Concept of Networking

At its most elementary level, a network consists of two computers connected to each other by a cable so that they can share data. All networking no matter how sophisticated, stems from that simple, system. While the idea of two computers connected by a cable may not seem extraordinary, in retrospect, it was a major achievement in communications.

Networking arose from the need to share data in a timely fashion. Personal computers are wonderful business tools for producing data spreadsheets, graphics, and others types of information, but do not allow you to quickly share the data you have produced, Without a network, the document have to be printed out so that others can edit them or use them. At best, you give files on floppy disks to others to copy to their computers. If others make changes to the document there is no way to merge the changes. This was, and still is, called working in stand along environment.




Figure 1.1 Simple Network

If the worker want to connect his computer to other computers, he could share the data on the other computers and the printers. A group of computers and other devices connected together is called a network, and the concept of connected computers sharing resources is called networking.

Computers that are part of a network can share the following:

- Data
- Message
- Graphics
- Printers
- Fax machines
- Modems
- Other hardware resources

This list is constantly growing as new ways are found to share and communicate by means of computers.

1.4 Benefits of Networking

Modern organization today is widely dispersed, with offices located in diverse parts of a country and the world. A network provides the means to exchange data among computers and makes programs and data available to the people.

Program and file sharing

Networking versions of many popular software packages are available at considerably lower prices when compared to buying individually licensed copies. The program and its data files are stored on a particular terminal for access by any network user.

Database sharing

A database program is an ideal application for a network. A network feature called record locking lets multiple users simultaneously access a file without corrupting the data. Record locking ensures that no two users edit the same record at the same time.

Sharing system resources

The system resources include printers, plotters, and storage devices, which can be shared by workstation, hooked on network. For example, consider the sharing of resource such as a storage devise. It is more efficient and economical to store an application on a network drive rather than storing a copy of it on

each of the user's local storage devices. Thus by reducing storage needs, hardware costs are also reduced. Similarly, a single printer (which is the most common shared resource) can suffice the requirement of multiple users on a network.

Back-up

Networking also provides the critical function of back-up. In the event, that one computer fails, its counterpart can assume its functions and worked. Back-up capability is especially useful in systems like air traffic control.

Creation of workgroups

Groups are important in networks. They can consist of users who work in a Department or who are assigned to a special project. With NetWare, you can assign users to groups. and then give each group access to special directories and resources not accessible by other users. This saves the trouble of assigning access to each individual user. Network workgroups facilitate new "flat" corporate structures in which people for diverse and remote departments belong to special group projects.

Centralized Management

Because NetWare uses dedicated servers, these servers can be grouped in one location, along with the shared resources attached to them, for easier management. Hardware upgrades, software backups system maintenance, and system protection are much easier to handle when these devices are in one location.

Security

Security starts with the login procedure to ensure that a user accesses the network using his or her own account. This account is tailored to give the user access only to authorized areas of the server and the network. Login restrictions can force a user to login at one specific station and only during a specific time frame.

Access to more than one operating system

Today's network operating system provides connectivity to systems running a variety of other operating system. This feature is called interoperability and TCP/IP is the protocol, which plays significant role in achieving this connectivity.

Enhancement of the corporate structure

Networks can change the structure of an organization and the way it is managed. Users who work in a specific department and for a specific manage no longer need to be in the

same physical area. Their offices can be located in areas where their expertise is most needed. The network ties them to their department managers and peers. This arrangement is useful for special projects in which individuals from different departments, such as research, production, and marketing need to work closely with each other. Computer networks and communication systems provide the rapid exchange of information residing on computers throughout a country.

1.5 Types of Networks

Networks can start small and grow with an organization. The classification of networks is done based on geographical area coverage. Here is the list of the different types of available networks.

- Local Area Network (LAN)
- Metropolitan Area Network (MAN)
- Wide Area Network (WAN)
- Enterprise Area Network (EAN)

Local Area Network

A LAN is a shared communication system to which many computers are attached. A LAN, as its name implies, is limited to a local area. This has to do more with the electrical characteristics of the medium than the fact that many early LANs were designed for departments, although the latter accurately describes a LAN as well. LANs have different topologies, the most common being the linear bus and the star configuration.

Metropolitan Area Network (MAN)

A MAN is a backbone network that spans a metropolitan area and may be regulated by local or state authorities. The Telephone Company, cable services, and other suppliers provide MAN services to companies that need to build networks that span public rights-of-way in metropolitan areas.

Wide Area Network

A Wide Area Network spans inter city, interstate, or international boundaries, the connectivity between the LANs or individual computer systems is achieved using the public or private links. Though the use of telephone networks is very common due to its availability and greater coverage, the other options such as leased lines and satellite links have also gained popularity due to their support for high-speed data transfer.

The WANs are usually owned by multiple organizations as they are formed by interconnecting a variety of system from multiple organization (as well as individuals). Internet is the best example of a WAN.

The WANs work on much lower speeds than LANs. Compared to the LAN data rates, which are in 10s or 100s of MBPS, the WAN data rates are usually in KBPS or one or two MBPS (max). The WAN links supporting a data transfer rate of 34-35 MBPS are also available, but are very expensive.

Enterprise Area Network

As enterprise network interconnects all the computer systems of an organization regardless of their operating systems, communication protocols, application software's, or geographic location. It may therefore be a LAN, MAN or WAN. An organization can connect its Sun SPARC networks, NetWare networks, Windows NT networks, OS/2 networks running a variety of protocols such as TCP/IP, NetBEUI, IPX/SPX etc, and different applications, to form an enterprise networks build maintain, and manage their own networks.

1.6 Classification of Network

Even with these similarities, networks can be divided into two broad categories:

- Peer-to-peer
- Server-based




Figure 1.2 Typical Peer-To-Peer and Server-Based networks

The distinction between peer-to-peer and server based networks is important because each has different capabilities. The type of network you implement will depend on numerous factors, including the:

- Size of the organization
- Level of security required
- Type of business
- Level of administrative available
- Amount of network traffic
- Needs of the network users
- Network budget.

1.6.1 Peer-to-peer-Networks

In a peer-to-peer network, there are no dedicated servers or hierarchy among the computers. All of the computers are equal and therefore are known as peers. Normally each computer functions as both a client and a server, and there is no one assigned to be an administrator responsible for the entire network. The user at each computer determines what data on their computer gets shared on the network.

Size

Peer-to-Peer networks are also called workgroups. The term workgroup implies a small group of people. In a peer-to-peer network, there are typically fewer than 10 computers.

Cost

Peer-to-peer networks are relatively simple. Because each computer functions as a client and a server, there is no need for a powerful central server, or for the other components needed for a high capacity network. Peer-to-peer networks can be less expensive than server based networks.

1.6.1.2 Peer-to-Peer Operating Systems

In a peer-to-peer network, the networking software does not need the same level of performance and security as the networking software designed for dedicated servers. Dedicated servers function only as servers and are not used as a client or workstation.

In operating systems such as Microsoft Windows NT Workstation, Microsoft Windows for Workgroups, and Microsoft Windows 95, peer-to-peer

networking is built into the operating system. No additional software is required to set up a peer-to-peer network.

Implementation

In a typical peer-to-peer environment there are number of networking issues that have standard solutions included.

- Computers located at the users' desks
- Users act as their own administrators and plan their own security
- A simple easily visible cabling system is used, which connects computer to computer in the network.

Where peer-to-peer is appropriate, Peer-to-Peer networks are good choices for environments where:

- There are fewer than 10 users
- The users are all located in the same general area.
- Security is not an issue.
- The organization and the network will have limited growth within the foreseeable future.

Considering these guidelines, there are times that a peer-to-peer network will be a better solution than a server-based network.

Peer-to-peer Consideration

While a peer-to-peer network may meet the needs of small organizations, this type of approach may be inappropriate in certain environments. The following networking areas illustrate some peer-to-peer issues, which a network planner will have to resolve before deciding upon which type of network to implement.

Administration

Network administration involves a variety of tasks including:

- Managing users and security.
- Making resources available.

- Maintaining application and data.
- Installing and upgrading application software.

In a typical peer-to-peer network there is no system manager who oversees administration for the entire network, each users administer their own computer.

Sharing Resources

All users can share any of their resources in any manner they choose. These resources include data in shared directories, printers, fax card, and so on.

Server Requirements

In a peer-to-peer environment each computer must:

- Use a large percentage of its resources to support the local user (the user at the computer).
- Use additional resources to support each remote user (a user accessing the server over the network) accessing its resources.

Security

Security consists of setting a password on a resource, such as a directory that is shared on the network. Because all peer-to-peer users set their own security and shares can exist on any computer rather only on a centralized server, centralized control is very difficult. This has a big impact on network security because some users may not implement any security at all. If security is an issue you should consider a server based network.

1.6.2 Server- Based Networks

In an environment with more than 10 users, a peer-to-peer network with computers acting as both servers and clients will probably not be adequate. Therefore, most networks have dedicated servers. A Dedicated Server is one that only function as a server and is not used as a client or workstation. Servers are “dedicated” because they are optimized to quickly service requests from network clients and to ensure the

security of files and directories. Server-based networks have become the standard model for networking.

As networks increase in size and traffic, more than one server on the network is needed. Spreading the tasks among several servers ensures that each task will be performed in the most efficient manner possible

Specialized Servers

The Variety of tasks that servers must perform is varied and complex. Servers for large networks have become specialized to accommodate the expanding needs of needs . For example, in a Windows NT Server network the different types of servers include the following:

File and print servers

File and print servers manage user access and use of file and printer resources. For example, if you were running a word processing application, the word processing application would run on your computer. The word processing document stored on the file and print server is loaded into your computer's memory so that you can edit or use it locally. In other words, file and print servers are for file and data storage.

Application servers

Application servers make the server side of client/server applications, as well as the data, available to clients. For example, servers store vast amounts of data that is structured to make it easy to retrieve. This differs from a file and print server. With a file and print server, the data or file is downloaded to the computer making the request. With an application server, the database stays on the server and only the results of a request is down loaded to the computer making the request.

A client application running locally would access the data on the application server. Instead of the entire data base being downloaded from the server to your local computer, only the results of your query would be loaded onto your computer. For example, you could search the employee database for all employees who were born in November.

Mail server

Mail servers manage electronic messaging between network users.

Fax servers


Fax servers manage fax traffic into and out of the network, by sharing one or more fax modem boards.

Communication servers

Communication servers handle data flow and e-mail messages between the server's own network and other networks, mainframe computers, or remote users using modems and telephone lines to dial in to the server.

Directory services server to enable users to locate, store, and secure information on the network. Windows NT Server combines computers into logical groupings, called domains, which allow any user on the network to be given access to any resources on the network.

Planning for various servers becomes important with an expanded network. The planner must take into account any anticipated network growths so that network use



will not be disrupted if the role of a specific server needs to be changed.

Figure 1.3 Specialized Servers

The Role of Software

A network server and the operating system work together as a unit. No matter how powerful or advanced a server might be, it is useless without an operating system that can take advantage of its physical resources. Certain advanced operating systems, such as Microsoft Windows NT Server, were designed to take advantage of the most advanced server hardware. For example Windows NT Server can take advantage of server hardware in the following ways.

Category	Feature
Symmetric MultiProcessing(SMP)	A multiprocessing system has more than one processor. SMP means that the system load and application needs are distributed evenly across all available processors
Multiple-Platform support	Faster processors from vendors such as Intel 386/486 and Pentium MIPS 4000 RISC and Digital Alpha AXP .
Filename/directory length	255 characters
File size	16 EB (264 bytes)
Partition Size	16 EB (264 bytes).

Server-Based Network Advantages

A server is designed to provide access to many files and printers while maintaining performance and security to the user.

Server based sharing of data can be centrally administered and controlled. The resources are usually centrally located and are easier to locate and support than resources on random computers.

For example, in Windows NT Server, directory resources are shared through File Manager.

To share a directory, highlight it and then from the Disk menu, select the Share as option.

Security

Security is most often primary reason for choosing a server based approach to networking. In a server based environment, such as a Windows NT server network , security can be managed by one administrator who sets the policy and applies it to every user on the network.

Backup

Because crucial data is centralized on one or a few servers, it is easier to make use that the data is backed up on a regular schedule.

Redundancy

Through redundancy systems, the data on any server can be duplicated and kept online. So that even if something happens to the primary data storage area for the data, a backup copy of the data can be used to retrieve the data.

Number of Users

A server based network can support thousand of users. This type of network would be impossible to manage as a peer-to-peer network, but current monitoring and network management utilities make it possible to operate a server based network for large numbers of users.

Hardware Considerations

Client computer hardware can be limited to the needs of the user because clients do not need the additional RAM and disk storage needed to provide server services. A typical client computer has at least a 486 processor and 8 to 16 MB of RAM.

1.7 Short Summary

- ☞ A computer network is a data communication system, which connects two or more computers and peripheral device.
- ☞ Network allows sharing resources, information and services.
- ☞ A network has both hardware and software components
- ☞ The hardware component include file server, workstation, NIC, cabling and resources like printers and tape devices
- ☞ A local area network can span a few kilometers(maximum)
- ☞ An enterprise network interconnects all the computer systems of an organization, regardless of their operating systems, communication protocols, application software's, or geographic location.

- ☞ The software components include network operating system such as server and workstation components
- ☞ The Classification of networks is done based on geographical area coverage.
- ☞ In a peer-to-peer network, there are no dedicated servers or hierarchy among the computers.
- ☞ Server based sharing of data can be centrally administered and controlled.

1.8 Brain Storm

1. What is Centralized Processing & Distributed processing?
2. What are the Benefits of Computer Networking?
3. Explain the concept of Networking?
4. Discuss about LAN.
5. What do you learn about MAN?
6. Give a notes on Peer-to-peer operating system?
7. What are the Various Types of Networking & How it is classified?
8. What are the pros of Server based Networks ?
9. Difference between peer to peer and server based networks?



Introduction to NT

Objectives

In this Lecture you will learn the following:

- Networking with Windows NT Server 4.0
- Knowing about Windows NT Server 4.0, the internet and Intranet
- Understanding the features Of NT

Coverage Plan

Lecture 2

- 2.1 Snap Shot
- 2.2 Networking with Windows NT Server 4.0
- 2.3 Windows NT Server 4.0, the Internet and Intranets
- 2.4 What New in Windows NT Server 4.0
- 2.5 Coming Attraction in Windows NT 5.0
- 2.6 Short Summary
- 2.7 Brain Storm

2.1 Snap Shot

Windows NT has become a very popular Operating System(OS) in a relatively short period of time. Since its debut in 1993, it has gone from being dismissed as a resource hogging OS that no one wanted to being the linchpin of Microsoft's Enterprise strategy and the focal point for Microsoft's BackOffice suite of enterprise level services.

This transformation is partly the result of Microsoft's willingness to keep developing and promoting NT despite its initially poor sales and partly the result of NT's continues technical growth. This growth has been matched by the expanding capabilities of PC hardware a high end system at NT's original ship date was 486/66, but Pentium, Pentium Pro, and RISC systems now provide enough computational power and I/O capacity to make full use of NT's features. This lecture describes what differentiates NT from other desktop and server operating systems.

- High performance, including support for native 32 bit applications, a faster file system, and full multitasking
- Support for large memory spaces, large disks, and multiple CPUs
- Security, auditing and recovery features
- Ability to integrate with NetBEUI, Macintosh, TCP/IP, Novell NetWare, and LAN Manager networks.

2.2 Networking with Windows NT Server 4.0

Networking is where Microsoft has made the greatest improvement in Windows NT. Microsoft claims that Windows NT Server 4.0's file services are more than twice as fast as the original version, and printing has been speeded up, too. Windows NT 4.0 includes Microsoft's new IPX/SPX stack, which appears to offer equal or better performance than Novell's own NetWare drivers. What's more important, however, is that Microsoft has adopted the Internet's venerable TCP/IP as the network protocol of choice, making Windows NT Server 4.0 more attractive to Microsoft's target market - Fortune 1000 firms. Unix servers running TCP/IP over 10 BaseT (10 Mbps unshielded twisted pair) ethernet cabling now dominate enterprise wide corporate local area networks. Although 60 percent or more of today's networked PCs may "speak" NetWare's IPX/SPX, by the end of the 1990s TCP/IP is likely to displace the standard Novell protocol in all but the smallest scale networks.

Windows NT 4.0 also thrives in heterogeneous networks by using a combination of TCP/IP, IPX/SPX and NetBEUI protocols. What's more you don't pay extra for Windows NT Server 4.0 capability to run simultaneous multiple network protocols.

Windows NT 3.1 Advanced Server established a new standard for ease of installation of a network operating system, and the setup program of Windows NT Server 4.0 from CD-ROM in about 30 minutes, and upgrade a Windows 95 or Windows for Workgroups 3.1 + peer-to-peer network with 20 to 30 clients in a day or so. On average, it takes about 15 minutes to reconnect each client to a Windows NT server 4.0 domain, including reconnecting clients to a relocated Microsoft Mail postoffice. You need a few more minutes per client if you use TCP/IP or IPX/SPX, rather than the Windows Network's simpler NetBEUI protocol.

Ease of installation - especially in a workgroup environment that might connect 20 to 100 clients - isn't the only benefit of using Windows NT Server 4.0. Since its inception, Windows NT Server has required substantially fewer administrative and support resources than its NetWare and UNIX competitors. Windows NT 4.0 offers various administrative tools, notably User Manager for Domains and Server Manager, with improved graphical user interfaces that simplify the life of network administrators. In the longer term, it is not the licenser fee and installation time that determines the economics of a network operating system - it's the annual administrative and support costs that make or break information system budgets.

2.3 Windows NT Server 4.0, the Internet, and Intranets

The remarkable growth of the Internet, brought about primarily by the proliferation of World Wide Web servers, is one of the principal contributors to increased adoption rate of Windows NT Server by organizations of all sizes. Today, most Internet servers run Unix, but Windows NT Server rapidly is gaining ground as the network operating system of choice for delivering Web pages. Windows NT's advantages cost less for the hardware and software needed to set up a Windows NT Web site, combined with easier administration and reduced support requirements than for Unix "boxes".

Microsoft arrived late at the Internet table, having waited until December 7, 1995 to elucidate its "Embrace and Extend" Internet strategy. A flurry of press releases and white papers announced Microsoft's intent to become a major player in the Internet server and browser markets. Subsequently, Microsoft released a torrent of free pre-beta (alpha or preview) and beta versions of Internet-related applications, programming tools, and add-ons. To gain market presence, Microsoft lets you download the latest versions of its Internet Information Server and Internet Explorer

browser from <http://www.micorsoft.com> for only the cost of connect time. Microsoft's objective in giving away these two products obviously is to increase the size of the market for Windows NT Server and Windows 95, respectively. Whether this strategy succeeds in displacing Netscape Navigator as the undisputed leader of the browser business remains to be seen. It's clear, however that much of the very rapid increase in sales of Windows NT Server during the first half of 1996 derived from the free IIS offer. You no longer need to download IIS and IE from Microsoft's Web site; IIS and IE are included on the Windows NT Server 4.0 CD-ROM.

The "real money " on the server side of the internet business comes from setting up private intranets, not creating internet sites. Intranets offer the convenience of allowing users to browse for information on a corporate Local Area Network or Wide Area Network by using a conventional Internet browser. Navigating hyperlinks to related HTML encoded documents with connections to server-resident applications is demonstrably easier for average PC users than running special-purpose, often complex client side applications. Conventional database query tools dedicated database front ends often require a substantial amount of user training. Inexperience users quickly gain a knack for finding the information they need by clicking text and iconic hyperlinks of web pages. Thus, organizations setting up Intranets minimize training costs and, because simple Web-based applications are relatively easy to code, save programming expense. Microsoft sells a license for Windows NT Sever with each free copy of IIS and gains the opportunity to sell a copy of Proxy Server to provide security for clients connecting to the Internet.

Microsoft announced its Site Server 2.0 and Site Server 2.0 Enterprise edition in May 1997. Site Server 2.0 is a \$1, 499 integrated package for deploying and managing Internet and Intranet Web sites, which includes the Personalization System, Content Replications System, Web Publishing Wizard, Posting Acceptor, Usage Analyst, and Site Analyst. Microsoft also includes "as a promotion" a copy of Visual InterDev for creating dynamic, database-related web content. The \$4,999 Enterprise Edition, designed for conducting electronic commerce on the web, adds a new version of Microsoft Commerce Server and upgraded versions of Usage Analyst and Site Analyst. The Enterprise Edition supports Microsoft Wallet , a payment system for Web-based purchases.

2.4 What's New in Windows NT Server 4.0

When the first rumors started to circulate about Cairo, the code name for the successor to NT 3.x quite a bit of speculation was about what it would include. Magazine writers and online pundits started calling Cairo "NT 4.0" as it was to be a major upgrade to 3x.

Since then, Cairo has moved from being a single release of software to a group of technologies that Microsoft will gradually introduce in to future releases of NTW and NTS. The first group of Cairo technologies actually appeared in Win 95, the revamped user interface (UI) and Plug and Play hardware support were originally slated for Cairo.

In the fall of 1995, Microsoft released a “preview” of NT 4.0 called Shall update Release (SUR). SUR basically added the Win 95-user interface to NT, but it was released as unsupported software, and many applications didn’t work with it. NT 4.0 includes a great deal more functionality than the SUR, some of which isn’t immediately apparent. Unless otherwise noted, all the changes discussed in the following text are in both NTW and NTS.

Windows 95 Interface

One change that’s immediately obvious is NT 4.0’s new look feel. Microsoft has adapted the Windows 95 user interface look and feel NT 4.0 incorporates the Win 95 shell components (the Start menu, Taskbar, desktop, and Explorer) as its own, including all the exported interfaces that other programs can use to extend the shell . The NT desktop also offers a revamped Task Manager and login manager.

NT 4.0 also includes the Win 95 common controls and dialog boxes. The common controls were actually included in NT 3.51, but most NT applications weren’t written to take advantage of them at the time. However many Windows 95 application can run on NT now that these controls are available. Finally, Version 4.0 of the Windows Help subsystem (which made its debut in Win 95) has made the transition as well. Some of the user interface changes, like the right-click context menu, rely on extensions made to the Object Linking and Embedding subsystem.

Enhancements to object linking and embedding (OLE)

For some time, Windows has supported the concept of embedding and linking data objects created by one application into another, this functionality, known as object linking and embedding OLE is the muscle behind the integration of Microsoft Office and a number of Office compatible application. OLE capable applications can share data while preserving the appearance and native content of the shared data.

Network OLE extends the ability to link and share information across the network . Instead of requiring that the data objects be copied when their container is moved. Network OLE supports cross network links. Microsoft stunned many observers by

announcing that Network OLE which most of the industry though would be delayed at least until 1997 would be included in NT 4.0.

Network OLE allows OLE clients and servers to call OLE reunites that are actually running on other machines on the network, instead of being limited to the same machine. For example, a large sever running an OLE capable CAD program could be an OLE server, serving CAD drawings to Microsoft Office users on network workstations. Users can share and modify data elements that reside permanently on the server . Network OLE transmits a representation of the data to the user's machine and allows the data to be edited remotely transparently to the user! Network OLE isn't the only new networking and communications feature in Windows NT 4.0,though as noted in the next section.

Networking and Communication

NT 4.0 includes some significant networking and communication capabilities that are new to the NT platform. The most promoted among these is Microsoft's Internet information server (IIS) a fast and well featured server for internet clients. Microsoft released it as freeware for NTS 3.51 in 1996, and it's included as part of NTS 4.0 as well.

NT 4.0 also includes full support for the Telephony API (TAPI), NT 3.51 supported a subset of TAPI, but full support means that NT 4.0 can run the Microsoft Exchange client and the new Dial-Up Networking subsystem – both of which are included. The Messaging API (MAPI), layer of NT 4.0 now supports the full set of extension added for communicating with Microsoft Exchange servers.

In previous NT releases, you should map TCP/IP addresses to host names, using the Windows Internet Name Service (WINS), but if you wanted to use the Internet standard Domain Name Service (DNS) protocol, you needed a third-party product. NT 4.0 includes a DNS service that integrates DNS with WINS to improve the mapping between computer names (semperfi), Internet address (semperfi.ingr.com), and TCP/IP addresses(129.136.283.14). NT 4.0 also improves its integration with NetWare networks. NTW users can now run NetWare login scripts and overall file and print communications with NetWare have been improved.

2.5 Coming attractions in Windows NT 5.0

Windows NT 5.0 will include all the enterprise edition add-ons described in the preceding section. Following are the most important publicly announced features of Windows 5.0

- ❖ Distributed File system (DFS) relies on a modified Jet Database, similar to that of Microsoft Exchange Server, to store file and folder information. Dfs lets you assign aliases to server shares so that shares of multiple servers appear to clients as having a single server name. Dfs also lets developers write “directory enabled” applications that can manipulate records in the Dfs databases.
- ❖ Kerberos 5.0 security uses session tickets to authenticate users, a process that’s faster than server authentication with Windows NT’s NTLM protocol. When you log on to Windows NT 4.0 with the pre-release version of Dfs installed, you have the option to use Kerberos security instead of Windows NT security.
- ❖ Active Directory Services (ADSI) uses the Internet-standard Lightweight Directory Access Protocol (LDAP), derived from the ITU’s X.500 directory and messaging standard. ADSI overcomes IT managers’ objections to Windows NT 4.0 domain-based directory system. Installing ADS under Windows NT 4.0 adds a Directories shortcut to you desktop that offers you the choice between viewing the conventional WinNT or more detailed LDAP directory structure. You use ASP with Internet Explorer 3+ to view and set the properties of directory entries.
- ❖ Microsoft Management Console (MMC) is a new system management tool container that houses Microsoft and third-party “Snap-ins” to administer network components, including Windows NT Servers.
- ❖ Windows Scripting Host (WSH) lets you execute batch files written in Visual Basic Script (VBScript or VBS) or JavaScript from Windows NT, either in Windows mode or from the command prompt.
- ❖ Win32 Driver Model (WDM) is a new driver architecture that lets hardware developers use the same 32 bit device drivers for Windows 9x and Windows NT 5.0
- ❖ 64-bit memory addressing lets Windows NT 5.0 take full advantage of 64 bit RISC processors, such as DEC’s Alpha series and the forth coming 64 bit Intel RISC processor, code-named Merced.

With the exception of WDM drivers and 64-bit memory addressing, the Windows NT server 5.0 features in this list are available as beta or final versions for use with version 4.0.

2.6 Short Summary

- Linear address space means programs can use memory without regard to what

physical segment it's in which greatly simplifies development.

- Objects created by one application into another, is known as object linking and embedding
- Network system supports a variety of protocols. Protocols are the set of network communication rules
- Both Windows NT Workstation and Windows NT Server provide 32-bit operating systems for fast, multitasking environments
- In Windows NT Workstation , Security features protect network and local resources and in Windows NT server, Security features include mandatory logon, discretionary access control, memory protection and auditing

2.7 Brain Storm

1. How Windows NT supports the native 32-Bit Environment?
2. In what way Windows NT will maintains its Security?
3. Explain the features of Windows 5.0 ?



Lecture 3

Understanding the Windows NT Operating System

Objectives

In this Lecture you will learn the following:

- Understanding the Features of Windows NT Operating System
- Overview of NT Architecture
- Shows how NT divides processing between user mode and kernel mode
- Understanding the concept of object oriented operating system

Coverage Plan

Lecture 3

3.1 Snap Shot

3.2 Windows NT Operating System Features

3.3 Windows NT System Architecture

3.4 Short Summary

3.5 Brain Storm

3.1 Snap Shot

Windows NT was designed to usurp UNIX's role in OS marketplace. This role is really two distinct roles. On one hand, NT was intended to be a desktop OS for high end users; on the other it was designed to be a robust, fast scalable server OS. To meet these objectives, NT was designed around a number of desirable characteristics

Windows NT is based on an executive that provides operating system services such as allocating memory, reading and writing disk files and controlling access to objects in the system. It's important to understand the two different modes in which Windows NT can run code: Kernel mode and User mode.

Kernel Mode

Most multi-user operating systems are based on the concept of a system kernel, a small piece of code that runs as privileged program and answers requests for services. Windows NT provides similar functionality in its Executive. The Windows NT Executive provides system services to other subsystems and applications. The Executive runs in kernel mode. Code executed in kernel mode has full access to all objects and hardware in the system. This unrestricted access is necessary to implement capabilities like the security monitor and virtual memory manager. Not all the Windows NT system functions run in kernel mode some are implemented in the protected subsystems.

User mode

By contrast with kernel mode, user-mode code has only limited access to system resources and services. The only way for a program running in user mode to receive a service like thread scheduling or memory allocation is to request it from the Executive.

3.2 Windows NT Operating System Features

- *Reliability* by protecting the core operating system from malfunctioning applications and by isolating the operating system and application from direct operations on hardware. Structured exception handling takes care of application processing and low-level errors. NTFS provides increased reliability for file operations by a built-in transaction logging system.

- *Extensibility* by adopting a client/server model that uses a base operating system extended by application programming interfaces. In this case, the term client/server is used in its single computer UNIX context, not that of the client/server model applied to networked database and internet/intranet applications.
- *Portability* across different processor platforms, including RISC systems, through the use of a processor-specific hardware abstraction layer that provides the isolation layer between the operating system and hardware. Portability of Windows NT is provided by writing the source code for the operating system, with a few exceptions, in an ANSI-standard C programming language.
- *Security* by compliance with at least the U.S. Department of Defense C2 standard, which provides “need-to-know” protection and auditing capability. Security in Windows NT primarily is implemented through ACLs.(Access Control List)
- *Compatibility* with existing 16 bit DOS and Windows applications, plus the most common PC hardware devices and peripherals. Windows NT also provides the capability to execute applications written to the POSIX.1 standard, a requirement of the federal government’s software procurement policies. Early versions of Windows NT supported NTFS, HPFS and FAT file systems. Windows NT 4.0 no longer handles HPFS volumes.
- *Scalability* for better performance through the use of multiple CPUs with a symmetrical multiprocessing(SMP) architecture. To take advantage of SMP, 32-bit Windows applications must be written to use multiple threads of execution.

3.3 Windows NT System Architecture

This lecture examines the specifics of NT’s system architecture, looking at how (and how well) NT achieve the goals set forth during its design, as well as how Microsoft’s methods of achieving some of those goals have changed during the product’s history.

In this discussion, we’re going to work from the bottom up. First, we’ll discuss some terminology and background for operating system architectures, then we’ll start at the hardware level and work up, talking about the hardware abstraction layer, the kernel, executive and user-mode services. Figure 3.1 contains an architectural diagram of Windows NT 4.0 and its various components.

.....

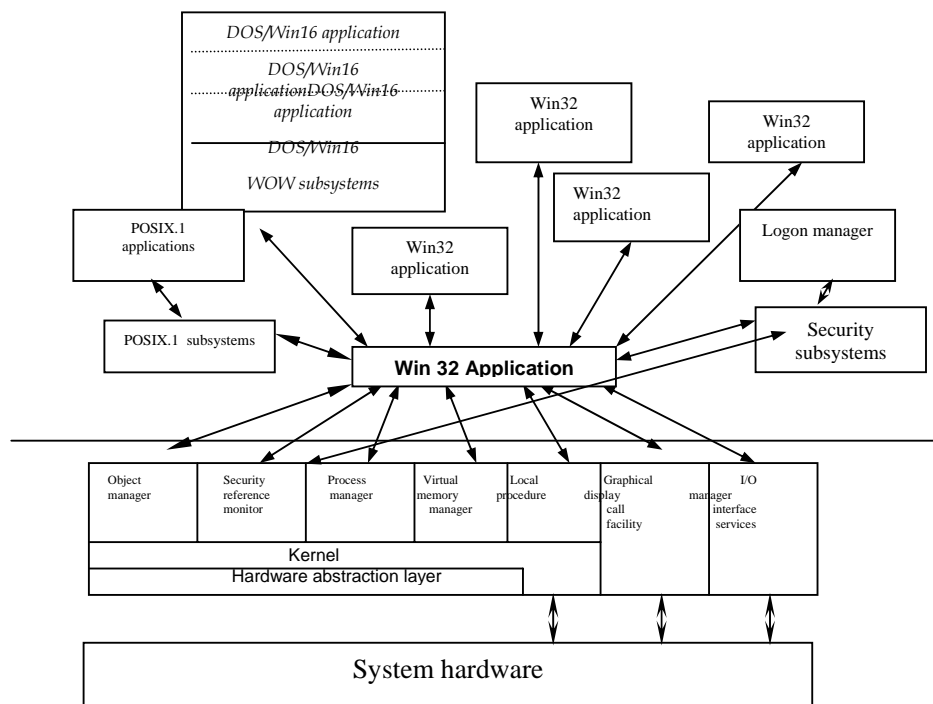


Figure 3.1 The Windows NT 4.0 Architecture

Let’s discuss NT’s various modules, starting from the bottom and working up. Many of NT’s modules interact with both the user and one another in order to accomplish their specified tasks.

I. The Hardware Abstraction Layer

The HAL component lies at the base of the NT architectural model. The HAL is coded in a mixture of C and processor specific assembly language. As you learned earlier, the HAL virtualizes access to many hardware functions, such as processor calls, memory and caches manipulation, and interrupt processing, giving a simple, virtual interface to the components that operate above it. This allows for the maximum degree of portability for operating system components in theory, only the HAL’s hardware specific code needs to be rewritten when porting to a new system architecture, with the rest of the operating system needing only to be recompiled on the new processor to complete the port.

As you can see in Figure 3.1, the HAL is not the only portion of the operating system that can access the hardware. The kernel, I/O manager, and graphical device interface services components all have some direct access to system hardware. These components all perform device-level functions that require some degree of hardware specific code are carefully isolated to minimize porting effort.

II. The Kernel

NT utilizes micro kernel architecture. This means that the kernel is responsible only for the scheduling and execution of tasks on system hardware. The kernel's software code cannot be preempted by any other portion of the operating system, the practice of paring down the kernel to a micro kernel, responsible only for task execution, enhances the stability and reliability of the operating system leaving only the smallest necessary piece of software with the ability to control all others.

The kernel creates, schedules, and closes threads on the system processors via the interface provided by the HAL. NT includes support for symmetric multiprocessing (SMP), which allows the kernel to divide threads among multiple system CPUs (if available), allowing for great increases in system performance by reducing the bottleneck of threads waiting for CPU execution. Threads can be scheduled with multiple priority levels, which determine the order of thread execution and (in the case of the real time priority) the ability of a thread to preempt lower priority threads.

III. Windows NT Executive Services

The NT Executive makes actual decisions about how resources are scheduled not the kernel the kernel merely carries out those system policies. This allow the kernel to remain static and not require rewriting if NT's designers revise the methods used for thread scheduling in later revisions of the operating system.

The Object Manager

Windows NT is made up of individual objects. NT takes the object-oriented model even further, by representing all operating system resources (such as files, processes, threads, memory segments, ports, and so on) as objects. These objects are created and handled by the Object Manager. The Object Manager creates each of these objects as requested by other operating system components, and grants these objects a standard set of interfaces and handles for their manipulation, this allow for generic creation, security protection and resource locking, client-use monitoring, and object resource tracking, and it provides a generic way to describe the object and its methods to other portions of the operating system.

The Security Reference Monitor

The NT Security Reference Monitor (SRM) implements security throughout the operating system on a per-object level. A request for creating, accessing, or destroying an object must pass through the SRM, which grants or denies access.

Each object in Windows NT has a security tag, known as its Access Control List (ACL). The ACL contains a list of NT security IDs of users and groups and the actions those users and groups are allowed to perform on each object.

The Security Reference Monitor works closely with the user mode Security Subsystem to generate and maintain security access tokens (SATs). When a user successfully logs on to an NT system, an SAT is generated containing the specific SID for that user as well as the SID's for all the groups he or she (or it, in the case of an automated system process).

The Security Reference Monitor uses these SATs to determine access to objects. When a user process attempts to manipulate a system object, the Monitor compares the SIDs in the user's SAT to the SID listed in the object's ACE and determines what, if any rights the user has to the specified objects.

The Process Manager

The Process Manager controls process objects. It is responsible for managing. The creation and deletion of process objects and the division of that process into threads, which can be manipulated by the system kernel. It provides APIs that allow applications to spawn and delete processes and split those processes into threads.

It does not schedule individual threads(that task is handled by the kernel)nor does it enforce policies as to process relationships (a task that is handled by individual environmental subsystems).

The Virtual Memory Manager

The Virtual Memory Manager handles memory access by system processes and the correspondence between virtual and actual memory addresses. Each process under NT can access a virtual address space of up to 4GB; of this space is used for the application's storage needs, and the other half is reserved for system use. The Virtual Memory Manager manipulated physical memory pages, contiguous chunks of system RAM that are 4KB. When an application requests data from a page in its address space, the Virtual Memory Manager determines whether the request page lies in RAM. If it does not, a page fault occurs, and the Virtual Memory Manager will retrieve that page from the system swap file. The Virtual Memory Manager also handles the process of moving unneeded pages out to the swap file and indexing and adjusting the size of the system swap file.

The Local Procedure Call Facility

The Local Procedure Call (LPC) facility provides NT applications and subsystems the ability to pass messages between one another, using named pipes. This signaling ability is crucial to NT's modular architectural design. LPC is similar to the Remote Procedure Call (RPC) facility, although LPC is optimized for use between elements within the operating system, while RPC is designed to be used over a network.

The Graphical Device Interface System

The graphical device interface (GDI) system handles access to graphics hardware by applications and environmental subsystems. It implements operating system graphic elements (such as windows, buttons, and list boxes) and their manipulation through the Window Manager. It also controls character-mode access to the screen, either in individual subsystem windows (such as POSIX or MS-DOS windows) or in full screen text mode windows. NT's GDI provides an API for 3D graphics design and manipulation using the open GL programming interface, first designed by Silicon Graphics. The GDI also implements DirectDraw, a subset of the new DirectX API which allows applications to perform faster graphic manipulation by giving them the ability to directly manipulate video hardware registers.

The movement of the GDI to the kernel mode Executive and the inclusion of DirectX APIs are new features to NT 4.0. Windows operating systems have traditionally provided poor video performance compared to that of MS-DOS, although the process of manipulating the screen is simpler under Windows than under DOS due to the APIs provided by the GDI, this abstracted process slows down the act of drawing objects on the screen. This made Windows unsuitable for game programming, which pushed video hardware to its limits. NT's video performance was worse still, due to its extreme isolation of system hardware from application software.

In order to facilitate the movement of the vast game market to Windows platforms Microsoft developed the DirectX API for Windows 95, which allowed programmers accelerated, direct hardware access to video (via DirectDraw) sound (using DirectSound) networking function (with DirectPlay) input

devices such as joysticks (with DirectInput) the playing Do onscreen movies (with ActiveMovie) and the generation and manipulation of 3D objects (using Direct3D) while still maintaining a generic case of APIs. The introduction of DirectX into Windows 95 has allowed the game market to port its development efforts easily to Windows 95, while simplifying the task of coding game software. With NT 4.0, Microsoft has now brought the Direct X API into NT, which should allow DirectX games and other application to run under NT 4.0 as well as Windows 95.

The movement of the GDI into the Executive also enhances video performance for everyday tasks, such as the user interface itself. NT 3.1 and 3.5x tended to suffer from slightly stilted graphics seeds, this was often made up for by NT's crisper multitasking, but when running multiple applications with many open windows, NT's video abstraction clearly affected interface performance. Now that the GDI runs in kernel mode, eliminating the need for constant processor context switching for each GDI call, the interface responds far more quickly and crisply than even previous versions of NT.

The I/O Manager

The I/O Manager handles all input/output function for NT. The I/O Manger uses driver software, customized for various devices, which can be swapped modularly in and out of the OS to customize it for different hardware uses. It also uses modular drivers for other services, such as file systems, which can be added removed and swapped to enhance the operating system's flexibility and reliability.

The I/O Manager handles all direct disk access, both for the system swap file (via the Cache Manager) and for individual files and directories (through file system and a -mass-storage device drivers.) It also controls network access (using network card drivers) and access to other system devices (such as sound cards, removable disk, and other hardware devotes.) NT 4.0 also includes the DirectX APIs, many of which are implemented within the I/O Manager. The IOM is comprised of these elements

The windows NT executive is a self contained, low level operating system, lacking only a user interface. The UI for the executive is provided by the win 32 subsystem that's discussed in the next section.

The IOM is unique among the executive components because it communicates with the kernel, with the HAL, and directly with hardware components through device drivers. The IOM is comprised of these elements :

- ❖ *File systems* support multiple file architectures through layered drivers thus, new or improved file systems can be added to Windows NT. Windows NT 4.0 supports 16 bit FAT and NTFS file systems OS/2's HPFS, which was supported in prior versions of windows NT can't be used with windows NT 4.0. for servers, the FAT file system is used only for accessing disks and CD-ROMs. (you can use the FAT file system for a windows NT server, but doing so is strongly discouraged.) Windows NT 4.0 does not support the 32 bit FAT file system included with service pack 2 for the OEM (original equipment manufacturer) version of Windows 95 that computer assemblers preinstall on their PCs. NTFS is the subject of the later section "Handling Files with NTFS.
- ❖ *Cache Manager* stores the most recently accessed fixed disk data in system memory to improve the performance of sequential reads. Cache Manager supplements the built in cache of high performance SCSI fixed disk and CD-ROM/WORM drives. Cache Manager also improves write performance by using VMM to perform asynchronous writes in the background.
- ❖ *Network drivers* consist of the network server and network redirector. The server receives network I/O requests, ad the redirector transmits network requests. The term redirector, which originated in the early days if MS-NET and PC-NET, is derived from the process of intercepting file operation requests on the local PC (interrupt int13h for Intel processors) and rerouting (redirecting)the operations over the network to the server.
- ❖ Device drivers consist of a 32 bit code layer between the preceding three elements sand hardware, such as SCSI host adapters for drives, network interface cards (NICs) keyboards, mice and graphic adapter cards. Windows NT treats network drivers as device drivers.

3.4 Short Summary

- ☞ The Executive runs in Kernel mode, its services allows subsystems to access system resources

- ☞ Object-oriented design and programming is based on a construct known as an object, developed to both simplify and enhance the process of writing complicated applications.
- ☞ An object is a unique set of programming code that performs a given task
- ☞ In the NT operating system HAL creates a bridge between the computer's operating system and hardware
- ☞ Kernel is the essential part of an operating system provides basic services
- ☞ The Graphical Device Interface (GDI) system handles access to graphics hardware by applications and environmental subsystems.

3.5 Brain Storm

1. Discuss the system architecture of Windows NT
2. What is an object oriented operating system?
3. Define the term 'Micro Kernel'?
4. How Kernel is used in Windows NT Architecture?
5. What is the role of HAL in Windows NT Architecture?
6. How Win32 application interact with NT Executive services?



Lecture 4

NT Environmental Sub Systems

Objectives

In this Lecture you will learn the following:

- Understanding the NT's Environmental subsystems
- Running NT application in Win32, WOW , OS/2 and POSIX.1 subsystems

Coverage Plan

Lecture 4

4.1 Snap shot

4.2 Client and Protected Subsystem Servers
--

4.3 Short summary

4.4 Brain Storm

4.1 Snap Shot

The layer above the Windows NT executive provide user-mode services. These layers consist of the environmental subsystems(servers) for 32-bit and 16-bit Windows applications, DOS, OS/2 and POSIX applications, plus client applications in each of these five categories. The Win32 subsystem provides the user interface for all servers and for those system services of the executive that involve interaction with users via for instances, User Manager for Domains, Server Manager and the Virtual Memory settings on the Performance page of the System Properties sheet. Thus, the Win32 subsystem is the only component of Windows NT that's visible to users.

4.2 Client and Protected Subsystem Servers

The following figure 4.1 shows the relationship of the environmental subsystems and client applications. Local procedure calls(solid line) provide all communication between subsystems and between subsystems and between clients and subsystems. Subsystems can request native services by a system trap, shown as dotted line in Figure 4.1. Each subsystem has its own protected memory space. Figure 4.1 shows these environmental subsystems and explain below :

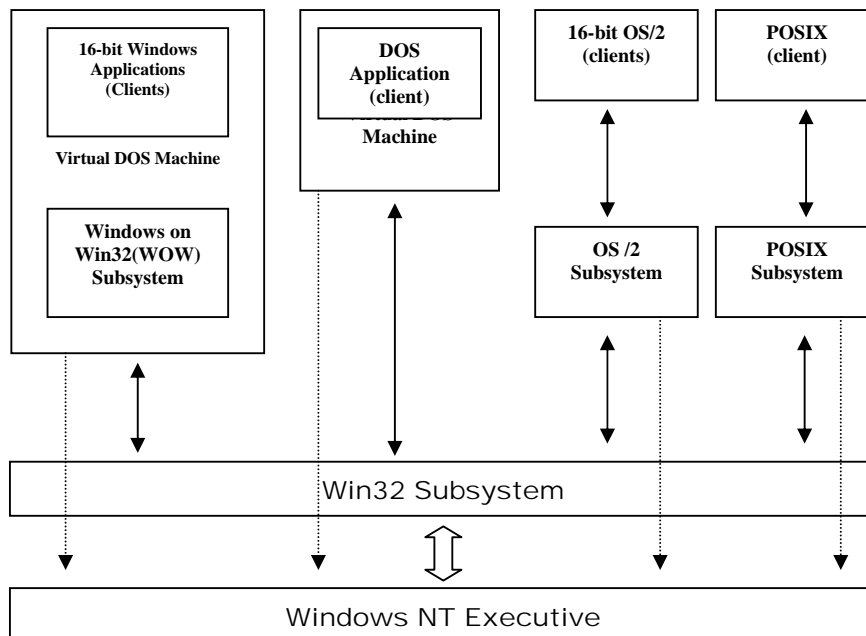


Figure 4.1 The environmental subsystems and client applications that provide user-mode services.

The Win 32 Subsystem

The Win32 subsystem, appropriately enough, implements the Win32 API. Win32 also manages keyboard and mouse inputs and handles display outputs for all other environmental subsystems. In fact, all the other subsystems running in user mode under NT are Win 32 applications, including the Security Subsystem, and the DOS/Win16 OS/2,POSIX subsystems.

Each Win32 application, whether a user program or a subsystem, has its own 4GB memory pool and system input queue, which processes system events such as keyboard and mouse manipulation and file system messages (see Figure 4.2). The use of independent input queues enhances NT's reliability, if a Win32 application crashes and stops fetch messages from its input queue, those messages do not block other applications from fetching their own messages from their own queues. Windows 3.1 applications on the other hand, share one message queue, if an application halts, its unretrieved messages can block the entire queue, preventing other application from processing their messages and hanging the entire operating system (or, in the case of 16 bit Windows applications running under Windows NT, the entire Win 16 subsystem) (see Figure 4.3).

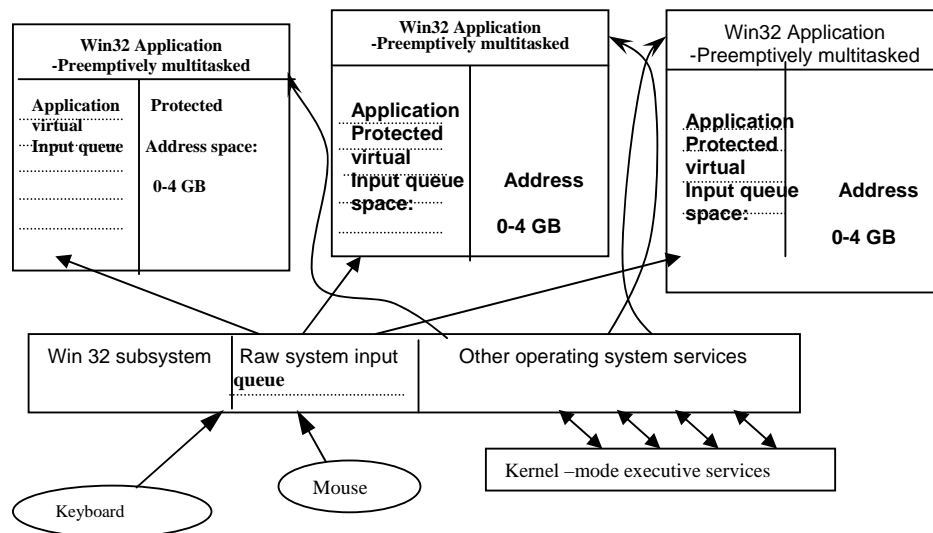


Figure 4.2 Win32 applications and the Win32 subsystems

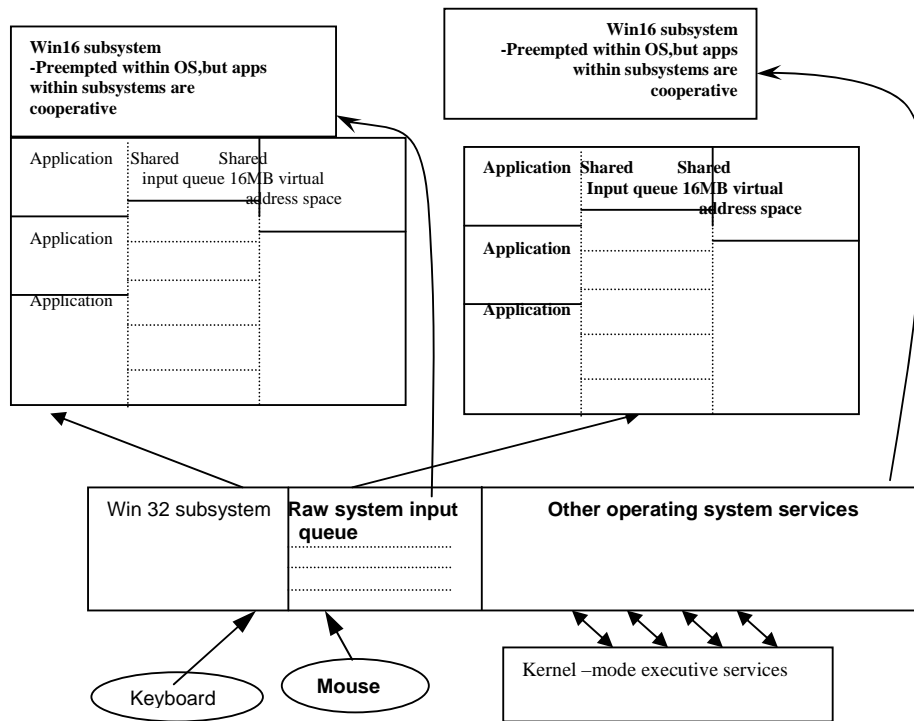


Figure 4.3 Win16 applications within the Win16 subsystems

The Win32 API is standard between different NT platform including the MIPS RISC system, DEC Alpha, Power PC, and of course the Inter x86 version of NT.

WOW Subsystem (Ms-Dos And Windows 16-Bit Environment)

The WOW subsystem provides a virtual DOS Machine (VDM) in the form of a Win32 application that emulated a 586 system (in NT 5.0) running MS-DOS and Windows 3.1. Each WOW subsystem can run as many DOS and Windows applications as its available resources can sustain, and since the WOW subsystem that can be run on an NT system.

When the user executes a DOS or Win 16 application, NT runs an instance of the WOW subsystem and executes the DOS/Win 16 code within it. When those applications terminate, the WOW subsystem remains active so that other 16 bit applications can be within it. NT provides the ability to create a new WOW subsystem upon execution of a 16-bit application, in order to provide it with memory protection from other 16 bit applications.

The VDM provided by the DOS subsystem allows most DOS application commands to be run within the WOW subsystem because it completely emulates an x86 processor running in Virtual-8086 mode. Some instructions, however, must be emulated by the Win 32 subsystem because they would violate NT's system integrity if allowed to function as normal on an MS-DOS system (in particular, input/output commands are emulated.) Some DOS applications do not respond well to this emulation, however, and may crash (or simply not run at all) under NT.

Windows 3.1x applications are also supported by the WOW subsystem, which provides translation between Windows 3.01x API calls and the Win 32 calls used natively by NT's Win32 subsystem. Again, some 16 bit Windows applications that use direct hardware access do not respond well to NT's emulation of their calls and may crash under NT.

Within each WOW subsystem application are cooperatively multitasked and share the same virtual memory pool, just as they would under Windows 3.1. Within a WOW subsystem, DOS and Windows applications can affect one another in precisely the same ways as on a native DOS/Windows system. Memory protection is not enforced, and multitasking is done cooperatively. However, as mentioned earlier the user can specify the creation of a new WOW subsystem when starting a 16 bit application, in order to offer that application protection from other 16 bit apps on the system.

Note that on other NT platforms there is no hardware support for executing x86 instructions natively, therefore, the Win16 subsystem on an Alpha, MIPS, or PowerPC platform requires the additional layer of a software emulator, which translates the 86x binary commands into native machine language. This imposes a significant performance hit on 16 bit applications.

OS/2 Subsystem

The OS/2 subsystem allows OS/2 1.x applications that do not use the OS/2 Presentation Manager GUI to run under NT Intel x86 systems. NT also includes the OS/2 HPFS file system to provide compatibility for OS/2 applications, as well as to ease the transition from OS/2 LAN Manager servers to Windows NT. The NT OS/2 subsystem does not support applications written for OS/2 2x, OS/2 Warp and Warp Connect, or OS/2 Merlin.

Non-Intel NT platforms do not include OS/2 subsystem support. Real mode OS/2 applications may run properly under the WOW subsystem, however.

Posix Subsystem

POSIX stands for Portable Computing System Interface for Computing Environments. It is a set of IEEE standards designed to provide a level of cross-OS compatibility for applications. The POSIX standard consists of a number of standards labeled POSIX.1, POSIX.2, POSIX 3., and so on. The standards ranges from basic source code compatibility to complete binary compatibility.

NT's POSIX subsystem supports only the POSIX .1 standard, which specifies source code compatibility for application written in the C language, using UNIX based API calls between applications and the operating system. Thus, any application that is written strictly to the POSIX.1, does not include such critical APIs as networking and security access, however, nearly any POSIX compliant application will require at least some receding to be compatible with NT.

NT includes support within the operating system and NTFS file system for POSIX requirements as well, such as the ability to disable traverse checking the ability to use case sensitive filenames, and other requirements.

4.3 Short Summary

- ✚ Win 32 subsystem does not provide binary compatibility between various platforms
- ✚ Win32 applications need to be recompiled for each intended release platform.
- ✚ WOW subsystems being a Win 32 applications is preemptively multitasked and each has a protected memory pool from other win32 application and subsystems
- ✚ No one applications can take the control of system resources or access the memory of other application
- ✚ If the user wants to executes a DOS or Win16 application, NT runs as instance of the WOW subsystem and the code within it
- ✚ Non-Intel NT platforms do not include OS/2 subsystem

✎ POSIX actually provide a level of cross – OS compatibility for applications

4.4 Brain Storm

1. Discuss about NT Environmental subsystems
2. How Win32 subsystems work in NT environment?
3. Whether OS/2 subsystem supports non- Intel Platforms? Explain?
4. What POSIX stands for?



Choosing Network Protocols

Objectives

In this Lecture you will learn the following:

- Understand network transport protocols
- About Functions of protocols
- Choosing protocols for client PC
- Knowing the advantages of multiple protocol

Coverage Plan

Lecture 5

- 5.1 Snapshot
- 5.2 Understanding the OSI Seven-Layer Model
- 5.3 Networking with Windows NT's Protocols
- 5.4 Supporting a Variety of PC Clients
- 5.5 Short Summary
- 5.6 Brain Storm

5.1 Snapshot

Protocols are rules and procedures for communication. For example, diplomats from one country adhere to protocol to guide them in interacting with diplomats from other countries. The use of communication rules applies in the same way in the computer environment. When several computers are networked, the rules and technical procedures governing their communication and interaction are called protocols.

There are three points to keep in mind when thinking about protocols in a network environment:

1. There are many protocols, while each protocol allows a basic communications, they have different purposes and accomplish different tasks. Each protocol has its own advantages and restrictions.
2. Some protocols work at various OSI layers. The layer at which a protocol works describes its function.

For example, a certain protocol works at the physical layer, meaning that the protocol at that layer ensures that the data packet passes through the network adapter card and out onto the network cable.

3. Several protocols may work together in what is known as protocol stack, or suite. Just as a network incorporated functions at every layer of the OSI model, different protocols also work together at different levels in a single protocol stack.. The levels in the protocol stack map or correspond to the layers of the OSI model. Taken together, the protocols describe the entire stack's functions and capabilities.

5.2 Understanding the OSI Seven-Layer Model

Communication over a network is a complex task. To simplify the task of discussing and building networks, the OSI seven layer network model was developed by the international standards organization, a branch of the United Nations headquarters in Geneva. OSI applies to virtually all computer-networking situations. At the heart of OSI is the diagram shown in figure 5.1. Each layer represents one of the seven different aspects of networking. Layer 1 the physical layer, is the most concrete, consisting of components that actually can be touched. On the other hand, layer 7 the application is the most abstract, consisting of high level software.

7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

Figure 5.1 *The Seven layers of the Open System Interconnection (OSI) model*

The OSI model doesn't correspond precisely to commercial networking software products, some of which span several OSI levels or break an OSI level into several tasks. Instead, the model provides a useful way of discussing the complex task of arranging communication between computers.

The Physical Layer

The physical layer sends bits over the wire or over another connection, such as a fiber-optic-cable or wireless connection, between computers. It deals with the electrical signals that represent the 0 (off) or 1 (on) state of a bit traveling over the network cabling. The decision to use a particular type of network interface card or a choice between twisted pair (10BaseT, 100BaseT) and coaxial cable (10Base2) is a decision about the physical layer, which is implemented in networking hardware.

The Data Link Layer

The data link layer deals with *frames*-groups of bits transmitted over the network. It relies on the physical layer to actually send the bits. The data link layer ensures that frames sent over the network are received and, if necessary resends them. Ethernet and Token Ring are examples of data link layers, and each has a different layout for its frames.

The IEEE networking model, established by the IEEE project 802 committee divides the data link layer into two sublayers: Logical link control(LLC) and Media Access Control(MAC) as shown in the figure 5.2. The MAC sublayer handles interaction with the physical layer below, and the LLC sublayer

handles interaction with the network layer above. The MAC layer provides standard interfaces for Carrier Sense Multiple Access with Collision Detection (CSMA/CD, Ethernet) networks, token-passing bus networks(similar to ARCnet,) and Token Ring networks.

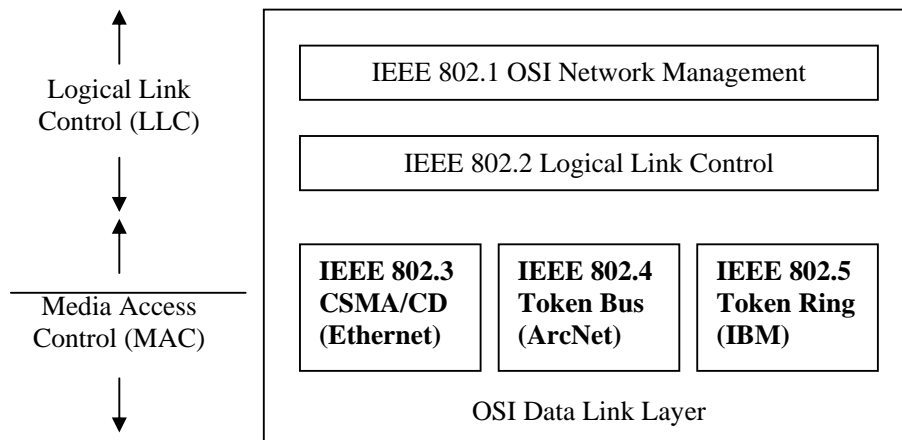


Figure 5.2 *The IEEE 802 series standards for Logical Link Control and Media Access Control within the OSI data-link layer.*

The Network Layer

The network layer deals with packets, which may be larger or smaller than frames. If the packets are larger than frames, the network layer breaks the packet into frames to send them, and reassembles them on receipt. If the packets are smaller than frames, the network layer bundles frames into packets to send them and breaks them apart on receipt.

In either case, the network layer relies on the data link layer to transmit the frames themselves. The network layer also deals with routing packets between computers on the network, and it knows the addresses of the hosts on the network. Typically the network layer can adjust the routing of packets to deal with network traffic and congestion. The network layer doesn't keep track of whether packets arrived at their destination or any errors occurred during transmission that job is handled by the transport layer.

The Transport Layer

The transport layer deals with message, which may be larger or smaller than packets. This ensures that messages travel between hosts without data loss

and, if necessary, arranges its packets to be re-sent. The transport layer relies on the network layer to transmit the frames.

NetBEUI, TCP/IP, IPX/SPX and other transport protocols handle the duties of the network layer and the transport layer. It is quite common for the network and transport layers to be combined into a single protocol.

The Session Layer

The session layer establishes and maintains a session between applications running on different computers. It knows the names of the computers on the network and handles security issues.

The session layer relies on the transport layer to transmit messages between the two computers. NetBIOS is the session layer for the Windows network, and sockets, the session layer for TCP/IP are examples of typical session layer software. Windows NT uses the 32-bit Windows sockets session layer.

The Presentation Layer

The presentation layer provides services that a number of different applications use, such as encryption, compression, or character translation (PC ASCII to IBM's EBCDIC, for example, or Intel's little-endian to Macintosh and Motorola's big-endian byte ordering.) The presentation layer relies on the session layer to pass on the encrypted, compressed, or translated material. The implementation of a presentation layer is XDR under RPC.

The Application Layer

The application layer handles requests by applications that require network communication such as accessing a database or delivering e-mail. This layer is directly accessible to applications, running on networked computers. It relies on the presentation layer to manipulate and transmit the communication. RPC is an example of an application layer implementation.

Many important protocols span the presentation and application layers for example named pipes and FTP (the file transfer protocol, not the application itself). Clients use named pipes, for example, to communicate with Microsoft SQL server. FTP is familiar to all UNIX and Internet users.

5.3 Networking with Windows NT's Protocols

Windows NT Server provides built-in support for seven networking protocols:

- TCP/IP
- NetBEUI and NetBEUI Frame (NBF)
- NW Link (IPX/SPX)
- Data Link Control (DLC)
- Apple Talk
- Remote Access Services (RAS)
- Streams

Figure 5.3 Shows how each protocol fits into the Windows NT network layer diagram . The following sections detail each of the seven protocols

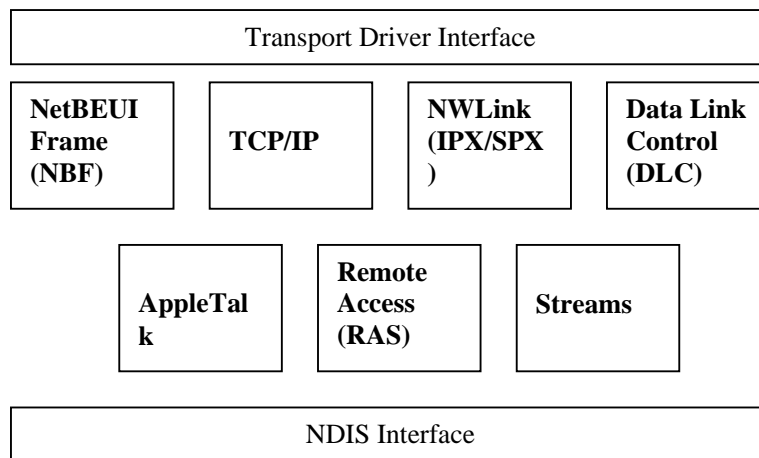


Figure 5.3 : Windows NT's transport protocols between the Transport Interface and the NDIS Interface

NetBEUI and NetBEUI Frame

NetBEUI is NetBIOS extended used interface. Originally NetBIOS and NetBEUI were very tightly tied together, and considered one protocol. However, several network vendors separated NetBIOS, the Session layer protocol, out so that it could be used with other routable transport protocols. NetBIOS (Network Basic Input/Output System) is an IBM Session layer LAN interface that acts as an application interface to the network. It provided the tools for a program to establish many application programs support it.

NetBEUI is a small, fast and efficient transport layer protocol that is supplied with all Microsoft network products. It has been available since the mid-1980s and was supplied with the first networking product from Microsoft, MS@NET.

NetBEUI networks are very simple to implement but are difficult to expand because the NetBEUI protocol isn't routable. NetBEUI advantages include its small stack size (important for MS-DOS-based computers,) its speed of data transfer on the network medium, and its compatibility with all Microsoft-based networks.

The NetBEUI Frame (NBF) protocol builds on NetBEUI to overcome some of these difficulties. In particular, NBF overcomes the session limit, which is 254 in NetBIOS. This makes it operable for larger networks that must retain compatibility with existing NetBEUI networks.

If you now have a NetBEUI network in operation, the commercial implementation of NetBEUI running on the PCs you plan to connect to your Windows NT server is likely one of the following:

- Microsoft LAN Manager, which uses a server that runs on Microsoft's implementation of OS/2
- IBM LAN Server, which uses a server running under IBM OS/2
- Microsoft MS-Net an early DOS peer - to - peer networking systems similar to IBM's original PC-Net

TCP/IP

Transmission Control Protocol/Internet Protocol (TCP/IP) is an industry standard suite of protocols providing communications in a heterogeneous environment. In addition, TCP/IP provides a routable, enterprise networking protocol and access to the worldwide Internet and its resources.

It has become the standard protocol used for interoperability among many different types of computers. This interoperability is one of the primary advantages to TCP/IP. Almost all networks support TCP/IP as a protocol. TCP/IP also supports routing, and is commonly used as an internetworking protocol. The size is not an issue and speed is about the same as IPX.

Because of its popularity, TCP/IP has become the de facto standard for Internet working. Other protocols written specifically for the TCP/IP suite include:

- SMTP (Simple Mail Transfer Protocol -E-mail
- FTP (File Transfer Protocol) For exchanging files among computers running TCP/IP
- SNMP (Simple Network Management Protocol) Network management

Historically, there were two primary disadvantages of TCP/IP its size and speed TCP/IP is a relatively large protocol stack which can cause problems in MS-DOS based clients. However, on Graphical User Interface (GUI) based operating systems, such as Windows NT or Windows, the size is not an issue and speed is about the same as IPX.

The most important advantage of TCP/IP over NetBEUI is that TCP/IP is a routable. A router is a device that forms a connection between a LAN and a WAN, or between two LANS. The router intercepts network packets and leaves them on the LAN if they're for a machine on a LAN; if not, the router passes the packets to another LAN or WAN. The structure of IP network addresses is designed specifically for efficient routing and the price of dedicated routing hardware has decreased dramatically over the past few years.

IP Addresses Every machine on a TCP/IP network has an IP address, such as 205.210.40.3. An IP address - sometimes referred to as a dotted quad - consists of four numbers, each in the range of 0-255, separated by dots. When an entire LAN joins the Internet IP addresses are commonly assigned to the LAN's machines that are easily distinguished from IP addresses on the rest of the Internet. Groups of related addresses are delineated into Class A, B, or C:

- ❖ A *Class C address* is actually about 250 IP addresses, each with the same values (for example, 205.210.40) for the first three components. The last component is different for each machine on the LAN. (A LAN with a class C address can't have 255 machines, because some values for each component are reserved.)
- ❖ A *Class B address* is actually about 60,000 IP addresses, each with the same values (for example, 130.105) for the first two components. Each of the last two components can vary. The owners of Class B addresses typically have far fewer than 60,000 machines on their internal networks, but more than 250.

- ❖ A *Class A address* is about 15 million IP addresses, all with the same first component (for example, 47) and with three different components at the end.

Only a limited number of these classes are available. For example, first components with values between 1 and 126 are reserved for Class A addresses (in practice, fewer than 50 Class A addresses have been assigned, primarily to Internet builders such as the U.S. military and telecommunications companies). First components between 128 and 191 are available for Class B addresses, and between 191 and 223 for Class C addresses. First components above 223 are reserved; if your installation requires only a single IP address, it's assigned from the Class B or C address of your ISP.

The IP address structure makes routing simple to arrange. If you have a Class C address, for example, traffic destined to an IP address that varies in the first three components—from that of your LAN is routed out to the rest of the Internet traffic with the same first three components as your LAN stays on the LAN. This routing is coded in a value called a subnet mask. The subnet mask 255.255.255.0 codes the routing pattern for a Class C address.

You can split a single Class C address into several smaller pieces, each piece on its own LAN. The subnet mask is bit-by-bit mask of the IP address, with each component representing 8 bits. For example, the subnet mask 255.255.255.224 defines a 31 address subnet with the following properties.

- ❖ All traffic for machines with the first three components varying from the routing machine is headed by the LAN.
- ❖ Of the remaining traffic, that for machines where the first 3 bits of the last component match those of the routing machine stays on the LAN. Because only the last 5 bits can vary, 31 ($16+8+4+2+1$) machines are on this subnet.
- ❖ The remaining traffic is headed by the LAN.

A Class C address could be split into seven such subnets allowing for great flexibility in combining small LANs. If large LANs are involved, a class B address can be split in a similar manner.

The Job of managing and assigning IP addresses and subnet masks for each client PC can involve substantial effort. The Dynamic Host Configuration Protocol (DHCP) included with Windows NT Server 4.0 reduces this effort significantly. To assign IP addresses by hand administrators must use Control Panel's Network tool; on the Protocols page, select the TCP/IP protocol and click Properties, choose the IP Address page, and then fill in the IP address and subnet mask

TCP/IP Printing TCP/IP provides support for printers connected to the network with an appropriate network interface card. IETF RFC 1179 defines the line printer (LPR) network protocol and utilities for TCP/IP printing. The line printer daemon (LPD) lets clients send print jobs to a spooler on a specified server, the server queues the spooled print job for the designated printer. Unlike TCP and IP, LPR and LPD are not formal Internet standards. Windows NT Server 4.0 provides several low-level enhancements to RFC 1179 TCP/IP printing.

Commercial TCP/IP Implementations : TCP/IP implementations are available for UNIX systems, the Macintosh, and Windows 3.x, and are included in Windows 95 and Windows NT 4.0 Server and workstation . A 32bit TCP/IP protocol stack for Windows 3.1+ is included in the Clients Tcp32 wfw folder of the Windows NT Server 4.0 distribution CD-ROM. TCP/IP supports many session -layer implementations. Including NetBT (NetBIOS over TCP/IP). Sockets (Winsock) , Streams, and RPC.

NWLink (IPX/SPX)

The Internetwork Packet Exchange / Sequenced Packet Exchange (IPX/SPX) protocol is the heart of Novel NetWare. Microsoft provides NWLink is its version of IPX/SPX. Like NetBEUI, It relatively small and fast protocol on a LAN. But, unlike NetBEUI, it does support routing. IPX /SPX is based on the Xerox Network System (XNS) protocol developed by Xerox Corp. XNS which is no longer in use, had two components: IDP (Internet Datagram Protocol) and SPP (Sequenced Packet Protocol). IPX is based on IDP; SPX is based on SPP.

IPX/SPX is a high performance protocol for LANs and is easier than TCP/IP to implement and administer. As with TCP/IP, IPX/SPX is a routable protocol, so it can be used to established a WAN.

Data Link Control (DLC)

IBM's Data Link Control (DLC) protocol is used to communicate with mainframes as a component of IBM's System Network Architecture (SNA). Windows NT 4.0 includes DLC interface device driver. On some LAN's DLC is used to communicate with printers connected directly to the LAN, rather than to a workstation or server. Several high-end Hewlett-Packard LaserJet printers offer DLC interfaces as an option. DLC printers are identified by a 12-byte numerical address for the network card.

Apple Talk

Apple Talk is Apple Computer's proprietary protocol stack designed to enable Apple Macintosh computers to share files and printers in a networked environment.

AppleTalk is the primary network protocol used by Macintosh computers, It's supported by Windows NT Server's Services for Macintosh, which allows Macintosh users to share Mac format files stored in Windows NT Server folders and use printers connected to a Windows NT server. Shared Windows NT folders appear to Mac users as conventional Mac folders Mac file names are converted to FAT (8.3) and NTFS standards, including long file names (LFNs) as required Dos and windows client applications that support Mac file formats can share the files with Mac users.

Windows NT Server supports the Mac file format, which consists of a resource fork and a data fork. Users can double-click a file stored on a Windows NT Server and by virtue of the signature resource of the resource fork, launch the associated application. The signature resource of the resource fork of a Mac file serves the same purpose as Windows file extension associations. Mac users can drag and drop their files from a Mac folder directly to a Windows NT server folder.

Remote Access Service (RAS)

Windows NT Server provides Remote Access Service (RAS, also called dial-in networking or Dial -Up Networking) to enable temporary connections to systems that aren't on your LAN—typically dial-up connections over a

conventional telephone line. Windows NT Server includes built-in support for Integrated Services Digital Network (ISDN) modems. A single Windows NT Server supports up to 255 simultaneous RAS connections.

The new PPTP (Point-to-Point Tunneling Protocol) – announced by Microsoft, 3Com Corp., Ascend Communications, ECI Telematics and U.S Robotics in March 1996 – allows users to “tunnel” through the Internet to reach their secure networks from a local dial-up connection.

Streams

Streams is a protocol specification that allows Windows NT to support third party communication protocols. In effect the third party Protocol is enclosed by an upper Streams layer and a lower Streams layer. The upper streams layer talks to the session layer and the third party protocol; the lower Streams layer talks to the third-party protocol and the data link layer.

Windows NT 3.1 used a Streams device driver for TCP/IP and IPX Streams has a substantial amount of overhead, so Microsoft abandoned the Streams protocol for TCP/IP and IPX/SPX effective with Windows NT 3.5 Windows NT 4.0 supports the Streams protocol for backward compatibility only .

5.4 Supporting a Variety of PC Clients

It is possible that you are installing an entire office full of computers, and you can choose the client PCs and operating systems that fit best with windows NT server. In this case, windows NT workstation 4.0 for desktop PCs and windows 95 for laptop PCs are the logical choices. It is more likely, however, that you are integrating a server into an environment that is now running one or more operating systems. The optimum protocol choice allows all your machines to talk to you server without spending large amounts of money or time.

It is possible for you server to support a number of different protocols at once, but to keep you workload to a minimum, choose you protocol(s) carefully so that you use as few as possible. In addition to making administration more complicated, running several protocols simultaneously consumes more memory on the server and client computers and can have a deleterious effect on

system performance. The following sections describe the types of clients commonly connected to windows NT servers and the networking protocols individual clients support.

Windows NT workstations

Not surprisingly, using clients running Windows NT workstation 4.0 allows you to use almost any of the networking protocols supported by Windows NT server including the following:

- ❖ NetBEUI Frame (NBF)
- ❖ NWLink (IPX/SPX)
- ❖ TCP/IP

In many cases, your choice of protocol under Windows NT is determined not by your operating system, but by the software your clients run. For example most windows Internet software runs over TCP/IP with Winsock, but not over NetBEUI or IPX/SPX.

PCs running Windows NT workstation 4.0 can participate in peer-to-peer networks, which operate independently of the Windows NT servers that support the primary (client/server) network. A maximum of 40 workgroup members can connect simultaneously to a Windows NT client that shares the workgroup files. NetBEUI is the most common method of establishing peer-to-peer workgroups. Most network administrators discourage peer-to-peer file sharing because of security issues, including lack of coordinated access control to sensitive information and centralized backup procedures.

Windows 95

Windows 95 offers the same suite of 32 bit network protocols as Windows NT workstation. One primary advantage of Windows 95 is the ease in which plug and play network adapter cards install. In general, software that runs under Windows NT workstation also can run under Windows 95 thus, users can choose their operating system without being concerned about networking issues. Another advantage of Windows 95 is that many new PCs are shipping with Windows 95 installed, saving the cost of another operating system.

Windows 95 has been designed for less powerful PCs but you should be aware of limitations (memory or disk space) on the client machines that preclude running Windows 95 efficiently.

Windows for workgroups 3.1x

Because it is an older product than Windows 95, windows for Workgroups (WfWg) 3.1x didn't support as many protocols when released. WfWg 3.1x limits the workstation to running only two networks at once. You connect to a Microsoft Windows network with NetBEUI and one other network by running one of the following add in protocols.

- IPX/SPX
- TCP/IP

Again the software you are running plays a major part in you choice among these options. Some software runs only over certain network protocols check you documentation.

Windows 3.1

Windows versions before windows for workgroups don't support networking directly; instead, early releases of windows used DOS-based networking. As with the DOS clients mentioned in the next section, Windows 3.1 clients attached to your Windows NT server can use one of the following protocols:

- NetBEUI
- IPX/SPX
- TCP/IP

It is not possible for one Windows 3.1 client to use more than one network protocol at a time.

DOS-only PCs

MS-DOS clients have a rather limited set of protocols available but can be connected to you Windows NT server by using one of these protocols:

- NetBEUI
- IPX/SPX
- TCP/IP

It is not possible for one MS-DOS PC to run more than one network protocol simultaneously. You use the drivers from the CD-ROM's \ Clients \ Msclient folder for DOS-only PCs.

OS/2 Clients

Just as Microsoft has improved PC networking capabilities with each new version of windows IBM has increased the networking features of OS/2, generally providing with the operating system access to protocol that in the past could be accessed only with third party add-ons. OS/2 versions 1.x 2.x and warp all can use any of these protocols.

- NetBEUI
- IPX/SPX
- TCP/IP

Apple Macintosh Computers

You can connect PCs that don't run DOS or Windows to a Windows NT server. The protocols supported for Macs are as follows.

- Apple Talk
- TCP/IP

Choosing apple Talk will make your difficult if any non Macintosh machines (other than you Windows NT server) are on your network and if you want the Macs to have peer-to-peer access to the files on non Macintosh workstations. If possible, all Macs connected to a network served by Windows NT should use the TCP/IP protocol.

Unix Workstations

Unix systems have even less in common with Windows and DOS systems than Mac systems do. However TCP/IP is a protocol that was designed for use with various operating systems. Using TCP/IP lets UNIX workstations talk to your windows NT server without difficulty.

5.5 Short Summary

- ☞ Protocols are rules and procedures for communicating.
- ☞ The Open Systems Interconnect model defines how transport protocols interact with a physical network and with the operating system.
- ☞ TCP/IP, NetBEUI, X.25, Xerox Network System (XNSTM), IPX/SPX and NW Link, APPC, Apple Talk, OSI protocol suite, DECnet are commonly used protocols.
- ☞ Window NT Server 4.0 makes it possible to use more than one network protocol simultaneously.
- ☞ Windows 95, Windows 3.1, Macintosh and UNIX clients can connect simultaneously to Windows NT Server 4.0
- ☞ Minimizing the number of protocols in use conserves resources and speeds network traffic.
- ☞ The groups of bits sent by the data-link layer is called a physical layer service data unit. In practice people call it a frame or a data.
- ☞ Winsock is the specification developed by more than 20 cooperating vendors that describes implementations of sockets on Windows machines.
- ☞ NetBEUI is a more efficient protocol for network browsing, especially over a Remote Access Service modem connection.

5.6 Brain Storm

1. What are the most commonly used Protocols?
2. Explain briefly about TCP/IP?
3. Explain the layers of OSI model?
4. Give a note on NetBEUI and NetBEUI Frame?



Lecture 6

Network Topologies and Architectures

Objectives

In this Lecture you will learn the following:

- Knowing about Access method
- Identify the three standard topologies and their varieties
- Understanding the concept of Network Architecture
- Advantages & Disadvantages of each topology

Coverage Plan

Lecture 6

6.1 Snap Shot

6.2 Access Methods

6.3 Network Topologies

6.4 Network Architecture

6.5 Short summary

6.6 Brain storm

6.1 Snap Shot

Networks come in a few standard forms or architectures and each form is a complete system of compatible hardware, protocols, transmission media and topologies. A topology is a map of the network. It is a plan for how the cabling will interconnect the nodes or devices and how the nodes will function in relation to one another. Several factors shape the various network topologies and one of the most important is the choice of an access method. An access method is a set of rules for sharing the transmission medium. This lecture describes two of the most important categories of access methods: contention and token passing. This lecture then looks at Ethernet, token-ring, ARCNet and FDDI networks. These types of networks all utilize either a connection-based or token passing access method.

6.2 Access Method

An access method is a set of rules governing how the network nodes share the transmission medium. The rules for sharing among computers are similar to the rules for sharing among humans in that they both boil down to a pair of fundamental philosophies:

- ◆ *Contention* in its purest form, contention means that the computers are contending for use of the transmission medium. Any computer in the network can transmit at any time (first come, first served).
- ◆ *Polling* one device is responsible for polling the other devices to see whether they are ready for the transmission or reception of data.
- ◆ *Token passing* The computers take turns using the transmission medium

As you can imagine, contention based access methods can give rise to situations in which two or more of the network nodes try to broadcast at the same time and the signals collide. Specifications for contention based access methods include procedures for how to avoid collisions and what to do if a collision occurs. This section introduces the CSMA/CD and CSMA/CA access methods.

On most contention based networks, the nodes are basically equal. No node has a higher priority than other nodes. A new access method called demand priority, however, resolves contention and collisions and in so doing accounts for data type priorities.

6.3 Network Topologies

A topology defines the arrangement of nodes, cables and connectivity devices that make up the network. Two categories form the basic for all discussions of topologies:

- ❖ Physical topology Describes the actual layout of the network transmission media
- ❖ Logical topology Describes the logical pathway a signal follows as it passes among the network nodes

Another way to think about this distinction is that a physical topology defines the way the network looks, and a logical topology defines the way the data passes among the nodes. At a glance this distinction may seem nit-picky, but as you will learn in this lecture, the physical and logical topologies for a network can be very different. A network with a star physical topology for example may actually have a bus or a ring logical topology.

In common usage, the word “topology” applies to a complete network definition, which includes the physical and logical topologies and also the specifications for elements such as the transmission medium.

Physical and logical topologies can take several forms. The most common and the most important for understanding the Ethernet and token-ring topologies are the following :

- ❖ Bus topologies
- ❖ Ring topologies
- ❖ Star topologies
- ❖ Mesh topologies

Bus Topology

The bus topology is also known as a linear bus. This is the simplest and most common method of networking computers. It consists of a single cable called a trunk (also backbone or segment) that connects all of the computers in the network in a single line.




Fig 6.1 Bus topology network

Communication of the Bus

Computers on bus topology network communicate by addressing data to a particular computer and putting that data on the cable in the form of electronic signals. To understand how computers communicate on a bus you need to be familiar with three concepts.

- Sending the signal
- Signal bounce
- The terminator.

Sending the Signal

Network data in the form of electronic signals is sent to all of the computers on the network however, the information is accepted only by the computer whose address matches the address encoded in the original signal. Only one computer at a time can send message.

Because only one computer at a time can send data on a bus network, network performance is affected by the number of computers attached to the bus. The more computers on a bus, the computers there will be waiting to put data on the bus, and the slower the network.

There is no standard measure for the impact of numbers of computers on any given network.. The amount of the network slows down the number of computers on the network. It depends on numerous factors including.

- Hardware capabilities of computers on the network
- Number of times computers on the network transmit data
- Type of cable used on the network
- Distance between computers on the networks.

The bus is a passive topology. Computers on a bus only listen for data being sent on the network. They are not responsible for moving data from one computer to the next. If one computer fails, it does not affect the rest of the network. In an active topology computers regenerate signals and move data along the network.

Signal Bounce

Because the data, or electronic signal is sent to the entire network, it will travel from one end of the cable to the other. If the signal were allowed to continue uninterrupted, it would keep bouncing back and forth along the cable and prevent other computers from sending. Therefore the signal must be stopped after it has had a chance to reach the proper destination address.

The Terminator

To stop the signal from bouncing, a component, called a terminator is placed at each end of the cable to absorb free signals. Absorbing the signal clears the cable so that other computers can send data

Every cable end on the network must be plugged into something. For example, a cable end could be plugged into a computer connector to extend the cable length. Any open cable ends- ends not plugged into something must be terminated to prevent signal bounce.

In bus topology, the workstation and the server are connected by one long cable. These cables form connectivity between the various computers in the network. The position of the server is not specific. Any workstation can be the file server provided it stores the network operating system to allow network users to share resources. The cable is also called the network bus and acts as a backbone to the network. Each workstation is connected to this bus using a shorter cable. The cable type is generally co-axial. This layout is quite popular in local area networks.

Some disadvantages are:

- If the cable fails then the entire system fails to respond to the use. Also since the cable is one, fault finding becomes very difficult.

Ring Topology

In this layout data flows in a circular manner through the ring cable. In most of such setups data only flows in one direction. The software to run this network is the simplest. The disadvantages are

- Data flows through workstations before reaching the server. If any one part of the cable between workstations or between workstation server is not working, the entire system will not work.




Figure 6.2 The Ring Topology

In other words, if the network is down the topology used for local area networks in an organization does not matter much. This is because the network covers a small geographic area. Hence, the efficiency of the network in terms of speed and the cost does not matter much. However, if the network is widespread, the information technology industry network users would opt for the star topology rather than the bus or the ring topology.

Star Topology

As the name specifies, this topology physically lays the workstation and the server in the form of a star. The server is at a central location whereas each workstation has an independent connection to the server. The independent connection of a workstation with the server is through a hub. A hub is similar to an extension board used for more than one connection to an electrical point. The hub has slots to connect various workstations with the server. Even if one cable is faulty, the other components continue to perform their tasks. Each connection is separate unlike a common one as in bus topology.

The software required to run this network is simpler. Each workstation interacts with the server independently. Hence, only data related to one workstation flows through the cable. Thus, the network operating system required to control such setups is simple. Figure 6.3 shows the star topology.

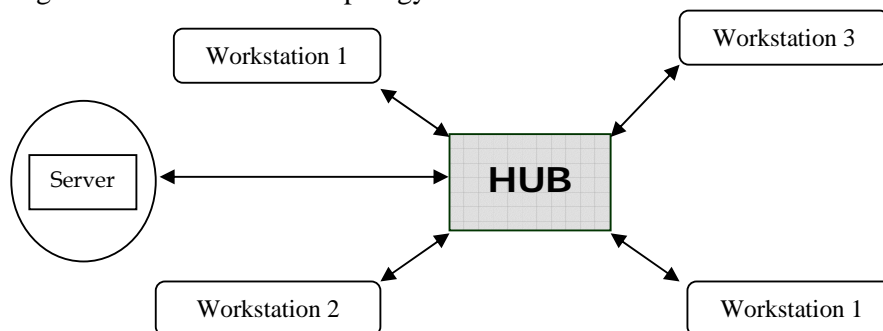


Figure 6.3 The Star Topology

Mesh Topology

A popular test subject is the mesh topology. A mesh topology is really a hybrid model representing an all-channel sort of physical topology. It is a hybrid because a mesh topology can incorporate all the topologies covered to this point. It is an all-channel topology, in that every device is directly connected to every other device on the network. When a new device is added, a connection to all existing devices must be made. This provides for a great deal of fault tolerance, but it involves extra work on the part of the network administrator. That is, if any transmission media breaks, the data transfer can take alternative routes. However, cabling becomes much more extensive and complicated. These different connection can be same (all Ethernet) or different (a mix of Ethernet and token ring) The below figure 6.4 shows the Mesh topology




Figure 6.4 *A mesh topology*

6.4 Network Architectures

A network architecture is the design specification of the physical layout of connected devices. This includes the cable being used (or wireless media being deployed), the types of network cards being deployed, and the mechanism through which data is sent on to the network and passed to each device. Network architecture, in short, encompasses the total design and layout of the network.

Ethernet

Ethernet is a very popular local area network architecture based on the CSMA/CD access method. The original Ethernet specification was the basis for the IEEE 802.3 specifications. In present usage, the term “Ethernet” refers to original Ethernet as well as the IEEE 802.3 standards. The different varieties of

Ethernet networks are commonly referred to as Ethernet topologies. Typically, Ethernet network can use a bus physical topology, although, as mentioned earlier, many varieties of Ethernet such as 10BASE-T use a star physical topology and a bus logical topology.

Ethernet networks, depending on the specification, operate at 10 or 100 Mbps using baseband transmission. Each IEEE 802.3 specification prescribes its own cable types.

The following are Ethernet topologies:

- ❖ 10BASE 2
- ❖ 10BASE 5
- ❖ 10BASE-T
- ❖ 10 BASE - FL
- ❖ 100 VG- Any LAN
- ❖ 100 BASE -X

Note that the name of each Ethernet topology begins with a number (10 or 100). That number specifies the transmission speed for the network. For instance, 10BASE5 is designed to operate at 10Mbps, and 100BASE-X operates at 100Mbps. "BASE" specifies that baseband transmission are being used. The "T" is for unshielded twisted pair wiring. "FL" is for fiber optic cable, "VG-AnyLAN implies voice Grade, and "x" implies media types.

Ethernet networks transmit data in small units called frames. The size of an Ethernet frame can be anywhere between 64 and 1,518 bytes, eighteen bytes of the total frame size are taken up by frame overhead, such as the source and destination addresses, protocol information and error checking information. There are many different types of Ethernet frames, such as the Ethernet II, 802.2, and 802.3 frame to name a few. It is important to remember that 802.2 and 802.3 are IEEE specifications on how information is transferred onto the transmission media as well as the specification on how the data should be packaged.

A typical Ethernet II frame has the following sections:

- ◆ Preamble. A field that signifies the beginning of the frame.

- ◆ Addresses. A field that identifies the source and destination addresses for the frame.
- ◆ Type. A field that designates the network layer protocol.
- ◆ Data. The data being transmitted.
- ◆ CRC. Cyclical Redundancy Check for error checking.

The term “Ethernet” commonly refers to original Ethernet which has been updated to Ethernet as well as the IEEE 802.3 standard. Ethernet and the 802.3 standards differ in ways significant enough to make standards incompatible in terms of packet formats, however. At the physical layer, Ethernet and 802.3 are generally compatible in terms of cables, connectors, and electronic devices.

Ethernet is generally used on light to medium traffic networks and performs best when a network’s data traffic transmits in short bursts. Ethernet is the most commonly used network standard.

One advantage of the linear bus topology used by most Ethernet networks (this doesn’t apply to star bus networks such as 10 BASE-T) is that the required cabling is minimized because a separate cable run to the hub for each node is not required, one disadvantage is that a break in the cable or a steaming network adapter card can bring down the entire network. Streaming is more frequently referred to as a broadcast storm. A broadcast storm when a network card fails and the transmitter floods the cable with traffic, like a faucet stuck open. At this point, the network becomes unusable.

Ethernet Cabling

You can use a variety of cables to implement Ethernet networks. Many of these cable types, such as Thinnet, Thicknet, UTP, and STP, are described. Ethernet networks traditionally have used coaxial cables of several different types. Fiber-optic cables now are frequently employed to extend the geographic range of Ethernet networks.

The contemporary interest in using twisted pair wiring has resulted in a scheme for cabling that uses unshielded twisted pair(UTP). The 10BASE-T cabling standard uses UTP in a star physical topology.

Ethernet remains closely associated with coaxial cable. Two types of coaxial cable still used in small and large environment are Thinnet (10BASE2) and Thicknet (10BASE5). Thinnet and Thicknet Ethernet networks have different limitations that are based on the Thinnet and Thicknet cable specifications. The best way to remember the requirements for Ethernet cable types is to use the 5-4-3 rule of thumb for each cable type.

The 5-4-3 rule states that the following can appear between any two nodes in the Ethernet network.

- Up to 5 segments in series
- Up to 4 concentrators or repeaters
- 3 segments of cable that contain nodes

Token Ring

Token ring uses a token passing architecture that adheres to the IEEE 802.5 standard, as described earlier. The topology is physically a star, but token ring uses a logical ring to pass the token from station to station. Each node must be attached to a concentrator called a multistation access unit.

In the earlier discussion of token passing, it may have occurred to you that if one computer crashes, the others will be left waiting forever for the token. MSAUs add fault tolerance to the network, so that a single failure doesn't stop the whole network. The MSAU can determine when the network adapter of a PC fails to transmit and can bypass it.

Token ring network interface cards can run at 4Mbps or 16 Mbps. Although 4Mbps cards can run at that data rate only, 16Mbps cards can be configured to run at 4 or 16 Mbps. All cards on a given network ring must run at the same rate. If all cards are not configured this way, either the machine connected to the card cannot have network access, or the entire network can be ground to a halt.

As each node acts as a repeater that receives tokens and data frames from its nearest active upstream neighbour. After the node processes a frame, the frame transmits downstream to the next attached node. Each token makes at least one trip around the entire ring and then returns to the originating node.

Workstations that indicate problem send a beacon to identify an address of the potential failure.

Token Ring Cabling

Traditional token ring networks use twisted pair cable. The following are standard IBM cable types for token ring:

- ✧ Type 1. A braided shield surrounds two twisted pairs of solid copper wire. Type 1 is used to connect terminals and distribution panels or to connect between different wiring closets that are located in the same building. Type 1 uses two STPs of solid core 22 AWG wire for long high data grade transmissions within the building's walls. The maximum cabling distance is 101 meters.
- ✧ Type 2. Type 2 uses a total of six twisted pairs; two are STPs and four are UTPs. This cable is used for the same purposes as Type 1, but enables both voice and data cables to be included in a single cable run. The maximum cabling distance is 100 meters (328 feet)
- ✧ Type 3. Used as an alternative to Type 1 and Type 2 cable due to its reduced cost, Type 3 has unshielded twisted pair copper with a minimum of two twists per inch. Type 3 has four UTPs of 22 or 24 AWG solid core wire for networks or telephone systems. Type 3 cannot be used for 16 Mbps token-ring networks. It is used primarily for low-data-grade transmission within walls. Signals don't travel as fast as with Type 1 cable because Type 3 doesn't have the shielding that Type 1 uses. The maximum cabling distance is 45 meters. Some vendors specify cabling distances of up to 150 meters.

Type 3 cabling (UTP) is the most popular transmission medium for token ring. A token-ring network using Type 3 (UTP) cabling can support up to 72 computers. A token-ring network using STP cabling can support up to 260 computers. The minimum distance between computers or between MSAUs is 2.5 meters.

A patch cable is a cable that connects MSAUs. Patch cables are typically IBM Type 6 cables that come in standard lengths of 8, 30, 75, or 150 feet. You can also get patch cables in custom lengths. You can use patch cables to extend the

length of Type 3 cables or to connect computers to MSAUs. Patch cables have an IBM connector at each end.

Token-ring adapter cables can have an IBM data connector at one end and a nine-pin connector at the other end, or they can use UTP cables with RJ-45 connectors on each end. Adapter cables connect client and server network adapters to other network components that use IBM data connectors. The type of connectors you need for a token-ring network depends on the type of cabling you're using. Type 3 cabling uses RJ-11 or RJ-45 connectors (Media filters, if necessary, can convert the network adapter to RJ-11 or RJ-45 format). Meanwhile, Type 1 and 2 cabling use IBM Type A Connectors.

Token-ring networks come in a few sizes and designs. A small movable token-ring system supports up to 12 MSAUs and uses Type 6 cable to attach clients and servers to IBM Model 8228 MSAUs. Type 6 is flexible but has limited distance capabilities. The characteristic of type 6 cable is suitable for small networks and for patch cords.

A large nonmoveable system supports up to 260 clients and file servers with up to 33 MSAUs. This network configuration uses IBM Type 1 or Type 2 cable. The large nonmovable system also involves other wiring needs such as punch panels or distribution panels, equipment racks for MSAUs and wiring closets to contain the previously listed components.

The MSAU is the central cabling component for IBM token-ring networks. The 8228 MSAU was the original wiring hub developed by IBM for its IBM Token-Ring networks. (IBM names all its hardware with numbers) Each 8228 has ten connectors, eight of which accept cables to clients or servers. The other connectors are labeled RI (ring in) and RO (ring out). The RI and RO connectors are used to connect multiple 8228s to form larger networks. The last RO must be connected to the first MSAU's RI.

8228S are mechanical devices that consist of relays and connectors. Their purpose is to switch clients in and out of the network. Each port is controlled by a relay powered by a voltage sent to the MSAU from the client. When an 8228 is first set up, each of these relays must be initialized with the setup tool that is shipped with the unit. Insert the setup tool into each port and hold it there until a light indicates that the port is properly initialized.

When you connect a token-ring network make sure you should do the following:

1. Initialize each port in the 8228 MSAU by using the setup tool shipped with the MSAU.
2. If you're using more than one MSAU, connect the RO port of each MSAU with the RI port of the next MSAU in the loop.
3. Connect the last RO with the first RI to complete the loop so that the MSAUs form a circle or ring.

Passing Data on Token Rings

As this lecture has already described, a frame called a token perpetually circulates around a token has control of the transmission medium. The actual process is as follows:

1. A computer in the ring, captures the token.
2. If the computer has data to transmit it holds the token and transmits a data frame. A token-ring data frame contains the fields.
3. Each computer in the ring checks to see whether it is the intended recipient of the frame.
4. When the frame reaches the destination address, the destination PC copies the frame to a receive buffer, updates the frame status field of the data frame, and puts the frame back on the ring.

In 16Mbps token-ring networks, the sending device can utilize an optional enhancement known as *early token release*. This is where the sending device issues a token immediately after sending a frame, not waiting for its own header to return. This speeds up the data transfers on the network.

5. When the computer that originally sent the frame receives it from the ring, it acknowledges a successful transmission, takes the frame off the ring it acknowledgement a successful transmission, takes the frame off the ring, and places the token back on the ring.

The Beaconsing Process

Generally, the first station that is powered up on a token-ring network automatically becomes what is called *active monitor* station. The responsibility of the active monitor station is to announce itself to the next active downstream station as the active monitor station and request that station to announce itself to its next active downstream station. The active monitor station sends this beacon announcement every seven seconds.

After each station announces itself to its next active downstream neighbour the announcing station becomes the nearest active upstream neighbor (NAUN) to the downstream station. Each station on a token-ring network has an upstream neighbor as well as a downstream neighbor.

After each station becomes aware of its NAUN, the beaconsing process continues every seven seconds beaconsed announcements from its upstream neighbor, it attempts to notify the network of the lack of contact from the upstream neighbor. It sends a message out onto the network ring, which includes the following:

- The sending station's network address
- The receiving NAUN's network address
- The beacon type.

From this information, the ring can determine which station might be having a problem and then attempt to fix the problem without disrupting the entire network. The process is known as autoreconfiguration. If autoreconfiguration proves unsuccessful, there may be a Ring Purge issued by the active monitor, forcing all computers to stop what they are doing and resynchronize with the ring. If both these mechanism fail manual correction becomes necessary.

ARC Net

ARCNet is an older architecture that is not found too often in the business world, but it have a presence in many older networks and school systems who often receive hand-me-downs from the business sector.

ARCNet utilizes a token-passing protocol that can have a star or bus physical topology. These segments can be connected with either active or passive hubs. ARCNet, when connected in a star topology, can use either twisted pair or coaxial cable (RG-62). If coaxial cable is used to create a star topology the ends of the cable can be attached directly to a BNC Connector, without a terminator. When in a bus topology, ARCNet uses a 93-ohm terminator which is attached to each end of the bus in a similar fashion to an Ethernet bus.

Each ARCNet card has a set of DIP switches built onto it. You can change the setting of these DIP switches to give each card a separate hardware address based upon these addresses. Tokens are passed to the card with the next highest address on the network. Due to this "access to the network passing" ARCNet shares some characteristics with a token passing network.

Some important facts about ARCNet are as follows:

- ARCNet uses a 93-ohm terminator (Ethernet uses a 50-ohm terminator)
- ARCNet uses a token-like passing architecture but does not require a NAU.
- The maximum length between a node and an active hub is 610 meters. The maximum length between a node and a passive hub is 30.5 meters
- The maximum network segment cable distance ARCNet supports is 6100 meters.
- ARCNet can have a total of only 255 stations per network segment.

FDDI

FDDI is very similar to token ring in that it relies on a node to have the token before it can use the network. It differs from token ring in that it utilizes fiber-optic cable as its transmission media allowing for transmissions of up to 100km. This standard permits up to 100 devices on the network with a maximum distance between stations of up to 2 kilometers.

FDDI has two different configurations: Class A and B. Class A uses two counteracting rings. Devices are attached to both rings. If one of these rings develops a fault, the other ring can still be used to transmit data. Class B uses a single ring to transmit data.

6.5 Short Summary

- ☞ Star topology, cables branch out from a hub
- ☞ In the ring topology each computer acts as a repeater and boosts the signal before sending it on.

- ☞ Hubs that regenerate and retransmit signals are active.
- ☞ The ring topology creates a single circle of cable, so there are no terminated ends.

6.6 Brain Storm

1. What are the different types of Topologies, explain briefly?
2. What are the disadvantages of Ring and Bus Topologies of Network Layout?
3. Explain the difference between a logical and a physical topology.
4. Explain the difference between contention-based, polling and token-passing access method.

☞☞☞

Lecture 7

Transmission Media

Objectives

In this Lecture you will learn the following:

- K About Network cables and types
- K Wireless Communications

Coverage Plan

Lecture 7

7.1 Snap Shot

7.2 Cable Media

7.3 Wireless Media

7.4 Comparisons of different Wireless Media

7.5 Short Summary

7.6 Brain Storm

7.1 Snap Shot

On any network the various entities must communicate through some form of media. Human communication requires some sort of media, whether it is technologically based(as are telephone wires) or whether it simply involves the use of our senses to detect sound waves propagating through the air. Likewise computers can communicate through cables, light and radio waves. Transmission media enable computers to send and receive messages but , as in human communication, do not guarantee that the messages will be understood.

This lecture discusses some of the most common network transmission media. One broad classification of this transmission media is known as bounded media or cable media. This includes cable types such as coaxial cable, twisted-pair cable and fiber-optic cable. Another media is known as boundless media: these media include all forms of wireless communication.

7.2 Cable Media

To connect different computers there is a line of communication. These lines of communication are called cables. These cables allow transferring of data and information from one computer to the other. They are the basic connecting links.

There are various kinds of cables that are used to connect the computers in the network. There are:

- ✧ Twisted Pair cables
- ✧ Co-axial cables
- ✧ Fiber Optic cables

Twisted-Pair Cable

Twisted-pair cable has become the dominant cable type of the new network designs that employ copper cable. Among the several reasons for the popularity of twisted-pair cable, the most significant is its low cost. Twisted-pair cable is inexpensive to install and offers the lowest cost per foot of any cable type. Your telephone cable is an example of a twisted pair cable.

A basic twisted-pair cable consists of two strands of copper wire twisted together. The twisting reduces the sensitivity of the cable to EMI (Electromagnetic Interference) and also reduces the tendency of the cable to radiate radio frequency noise that interferes with nearby cables and electronic components, because the radiated signals from the twisted wires tend to cancel each other out.(Antennas, which are purposely designed to radiate radio frequency signals, consist of parallel not twisted wires)

Twisting of the wires also controls the tendency of the wires in the pair to cause EMI in each other. As noted previously whenever two wires are in close proximity, the signals in each wire tends to produce cross talk in the other. Twisting the wires in the pair reduces crosstalk in much the same way that twisting reduces the tendency of the wires to radiate EMI. Two types of twisted-pair cable are used in LANs: Unshielded twisted-pair and Shielded twisted pair.




Figure 7.1 Unshielded twisted pair and Shielded Twisted Pair cables

1. Unshielded Twisted-pair(UTP)

UTP using the 10BaseT specification is the most popular type of twisted-pair cable and is fast becoming the most popular LAN cabling. The maximum cable length segment is 100 meters or about 328 feet.

UTP consists of two insulated copper wires. Depending on the particular purpose, there are UTP specifications, which govern how many twists are permitted per foot of cable. In the North American continent, UTP cable is the most commonly used cable for the existing telephone systems and is already installed in many office buildings.

UTP is specified in the Electronic Industries Association and the telecommunications Industries Association 568 commercial buildings wiring standard. EIA/TIA 568 used in creating standards that apply to a variety of building and wiring situations and ensure consistency of products for customers. These standards include five categories of UTP:

- *Category 1*
This refers to traditional UTP telephone cable which can carry voice but not data. Most telephone cable prior to 1983 was Category 1 cable.
- *Category 2*
This category certifies UTP cable for data transmissions up to 4 Mbps (megabits per second.) It consists of four twisted pairs.
- *Category 3*
This category certifies UTP cable for data transmissions up to 10 Mbps. It consists of four twisted pairs with three twists per foot.

- *Category 4*

This category certifies UTP cable for data transmissions up to 16 Mbps. It consists of four twisted pairs.

- *Category 5*

This category certifies UTP cable for data transmissions up to 100 Mbps. It consists of four twisted pair of copper wire.

Most telephone system use a type of UTP. In fact, one reason why UTP is so popular is because many buildings are prewired for twisted pair telephone systems. As part of this prewriting extra UTP is often installed to meet future cabling needs. If preinstalled twisted pair cable is of sufficient grade to support data transmission, it can be used in a compute network. Caution is required, however, because common telephone wire may not have the twisting and other electrical characteristics required for clean, secure, computer data transmission.

One potential problem with all types of cabling is crosstalk. You may remember that crosstalk is defined as signals from one line getting mixed with signals from another line. UTP is particularly susceptible to crosstalk. Shielding is used to reduce crosstalk.

UTP cable is the least costly of any cable type, although properly installed Category 5 tends to be fairly expensive. Distance limits for voice cabling are much less stringent than for data-grade cabling.

Connectors for UTP

The most common connector used with UTP cables is the RJ-45 connector. These connectors are easy to install on cables and are also extremely easy to connect and disconnect. An RJ-45 connector has eight pins and looks like a common RJ-11 telephone connector. They are slightly different size, however and won't fit together: an RJ-11 has only four pins.

Distribution racks, trays shelves and patch panels are available for large UTP installations. These accessories enable you to organize network cabling and also provide a central spot for expansion and reconfiguration. One necessary accessory, a jack coupler is a small device that attaches to a wall plate or a patch panel and receives an RJ-45 connection. Jack couplers can support transmission speeds of up to 100 Mbps.

2. Shielded Twisted Pair

STP uses a woven copper braid jacket, which is higher quality more protective jacket than UTP, has. STP also uses a foil wrap between and around the wire pairs, and internal twisting of the pairs. This gives STP excellent insulation to protect the transmitted data from outside interface.

What this means is that STP is less susceptible to electrical interference and supports higher transmission rates over longer distances than UTP.

The twisted pair cables are used for telephone connections. They contain two wires of conducting material insulated from each other. They appear twisted hence the name. They are very cheap and easily available. They are nowadays rarely used for computer network interconnections to it. Hence data transmission is not error free.

STP cables costs is more than thin coaxial or unshielded twisted-pair cable. STP is less costly, however than thick coaxial or fiber-optic cable.

Connectors for STP

AppleTalk and Token Ring networks can be cabled using UTP cable and RJ-45 connectors, but both networks originated as STP cabling systems. For STP cable, AppleTalk also employs a DIN-type connector. An IBM Data Connector connected to a network card having a DIN (DB-9) connector using a STP cable.

The IBM Data Connector is unusual because it doesn't come in two-gender configuration. Instead any IBM Data Connector can be snapped to any other IBM Data Connector.

Coaxial Cable

At one time, coaxial cable was the most widely used network cabling. There were a couple of reasons for coaxial's wide usage. Coaxial was relatively inexpensive and it was light flexible and easy to work with. It was so popular that it became a safe, easily supported installation.

In its simplest form, coaxial consists of a core made of solid copper surrounded by insulation, a braided metal shielding, and an outer. One layer of foil installation and one layer of braided metal shielding is referred to as dual

shielded. However Quad shielding is available for environment that are subject to higher interference. Quad shielding consists of two layers of foil insulation and two layers of braided metal shielding




Figure 7.2 Coaxial cable showing various layers

The components of a coaxial cable are as follows:

- A center conductor (conducting core) , although usually solid copper wires, is sometimes made of stranded wire.
- An outer conductor (Copper wire mesh) forms a tube surrounding the center conductor. This conductor can consist of braided wires, metallic foil or both. The outer conductor frequently called the shield, serves as a ground and also protects the inner conductor from EMI.
- An insulation layer (insulation) keeps the outer conductor spaced evenly from the inner conductor
- The jacket (outer shield) protects the cable from damage.

Shielding refers to the woven or stranded metal mesh(or other material) that surrounds some types of cabling . Shielding protects transmitted data by absorbing stray electronic signals, called noise, so that they do not get onto the cable and distort the data. The core of a coaxial cable carries the electronic signals which make up the data. This core wire can be either solid or stranded. If the core is solid it is usually copper.

The core is surrounded by an insulating layer which separates it form the wire mesh. The braided wire mesh acts as a ground and protects the core from electrical noise and crosstalk. Crosstalk is signal overflow from an adjacent wire.

The conducting core and the wire mesh must always be separated from each other. If they touch, the cable will experience a short, and noise or stray signals on the mesh will flow onto the copper wire. This will destroy the data.

Coaxial cable is more resistant to interference and attenuation than twisted pair cabling. Attenuation is the loss of signal strength, which begins to occur as the signal travels further along copper cable.

The stranded, protective sleeve can absorb stray electronic signals so they do not affect data being sent over the inner copper cable. For this reason, coaxial is a good choice for longer distances and for reliably supporting higher data rates with less sophisticated equipment. The co-axial cables are most commonly used. This is the type of cable that also connects your television to the cable operator's setup.

The coaxial cable used for Thinnet falls at the low end of the cost spectrum, whereas Thicknet is among the more costly options.

Types of Coaxial Cable

The two basic classifications for coaxial cable are as follows:

- Thinnet (thin)
- Thicknet (thick)

What type you select depends on the needs of your particular network.

Thinnet

Thinnet is a light and flexible coaxial cable .25 inch thick. Because this type of coaxial is flexible and easy to work with it can be used in almost any type of network installation. Networks that use Thinnet have the cable connected directly to a computer's network adapter card. Thinnet cable can reliably transmit a signal for 185 meters (about 610 feet).

Thicknet

Thicknet (big surprise) is thicker than Thinnet. Thicknet coaxial cable is approximately 0.5 inches (13 mm) in diameter. Because it is thicker and does not bend as readily as Thinnet, Thicknet cable is harder to work with. A thicker

center core, however means that Thicknet can carry more signals a longer distance than Thinnet. Thicknet can transmit a signal approximately 500 meters (1,650 feet).

Thicknet cable is sometimes called Standard Ethernet (although other cabling types described in this lecture are used for Ethernet also). Thicknet can be used to connect two or more small Thinnet LANs into a larger network.

Because of its greater size, Thicknet is also more expensive than Thinnet. However, Thicknet can be installed relatively safely outside, running from building to building.

Connectors for Coaxial Cable

Two types of connectors are commonly used with coaxial cable. The most common is the British Naval Connector (BNC). Key issues involving Thinnet cabling are

- A BNC T-connector connects the network board in the PC to the network. The T-connector attaches directly to the network board.
- BNC cable connectors attach cable segments to the T-connectors
- A BNC barrel connector connects to Thinnet cables
- Both ends of the cable must be terminated. A BNC terminator is a special connector that includes a resistor that is carefully matched to the characteristics of the cable system.
- One of the terminators must be grounded. A wire from the connector is attached to a grounded point, such as the center screw of a grounded electrical outlet.

In contrast, Thicknet uses N-connectors, which screw on rather than use a twist lock. As with Thinnet both ends of the cable must be terminated and one end must be grounded.

Workstations don't connect directly to the cable with Thicknet. Instead, a connecting device called a transceiver is attached to the Thicknet cable. This transceiver has a port for an AUI connector (which looks deceptively like joystick connector), and an AUI cable (also called a transceiver cable or a drop cable) connects the workstation to the

Thicknet medium. Transceiver can connect to Thicknet cables in the following two ways:

- Transceivers can be connected by cutting the cable and splicing N-connectors and a T-connector on the transceiver. Because it is so labor-intensive, this original method of connecting is used rather infrequently.
- The more common approach is to use a clamp-on transceiver, which has pins that penetrate the cable without the need for cutting it. Because clamp-on transceivers force sharp teeth into the cable, they frequently are referred to as vampire taps. You can use a transceiver to connect a Thinnet LAN to a Thicknet backbone.

Fiber Optic Cables

The fiber optic cables carry data and information and transmit them in the form of light rays. Light sources include lasers and light-emitting diodes(LEDs). LEDs are inexpensive but produce a fairly poor quality of light suitable for only less-stringent applications. The end of the cable that receives the light signal must convert the signal back to an electrical form. Several types of solid-state components can perform this service. Hence, they are free from any outside interference like noise.

The cable consists of an internal light conducting material that is glass based, covered by a tough insulating material. They are very widely used in computer networks although they are costly. However, it is capable of transmitting data in huge volumes and high speeds. The chances of data being corrupted through the transmitting line are very less.

Fiber Optic Composition

A fiber-optic network cable consists of two strands separately enclosed in plastic sheaths. One strand sends and the other receives. Two types of cable configuration are available: loose and tight configurations. Loose configurations incorporate a space between the fiber sheath and the outer plastic encasements; this space is filled with a gel or other material. Tight configurations contain strength wires between the conductor and the outer plastic encasement.

Optical fibers consists of an extremely thin cylinder thin cylinder of glass, called the core, surrounded by a concentric layer of glass, known as the cladding. The fibers are sometimes made of plastic. Plastic is easier to install, but cannot carry the light pulses as far as glass.

Each glass strand passes signals in only one direction, so a cable consists of two strands in separate jackets. One strand transmits and one receives. A reinforcing layer of plastic surrounds each glass strand while kevlar fibers provide strength. See Figure 7.3 for an example of Kevlar fibers. The kevlar fibers in the fiber optic connector are placed between the two cables, which are encased in plastic.




Figure 7.3 *Fiber-Optic Cable*

Fiber optic cable transmission are not subject to electrical interference and are extremely fast (currently about 100 Mbps with demonstrated rates of up to 200,000 Mbps). They can carry signal the light pulse for miles.

Fiber-Optic Considerations

Use fiber optic cable if you:

- Need to transmit data at very high speeds over long distances in a very secure media.
- Do not use fiber optic if you are under a tight budget.
- Do not have the expertise available to properly install it and connect devices to it.

7.3 Wireless Media

The wireless environment is emerging as a viable networking option. As the technology matures, vendors will be offering more products at attractive prices, which, in turn, will mean increased sales and demand. As demand increases the wireless environment will grow and improve.

The phrase wireless environment is misleading because it implies a network completely free of cabling. In most cases, this is not true. Most wireless networks actually consist of wireless components communicating with a network that uses cables in a mixed component network called a hybrid.

Wireless capabilities

The idea of wireless networks is attracting attention because wireless components can:

- Provide temporary connections to an existing, cabled network
- Help provide backup to an existing network
- Provide a certain degree of portability
- Extend networks beyond the limits of copper or even fiber-optic cables.

Reasons for Wireless Networks

Difficulty implementing cable is a factor, which will continue to push wireless environments toward greater acceptance. Wireless can be especially useful for networking.

- Spaces where cabling would be impossible or inconvenient, These include open lobbies, inaccessible parts of buildings, older building, historical building where renovation is prohibited and outdoor installations.
- People who moved around a lot within their work environment. Network administrators, for instance must troubleshoot a large office network. Doctors and nurses need to make rounds at a hospital.
- Temporary installations. These situations include any temporary department set up for a specific purpose that soon will be torn down or relocated.
- People who travel outside of the work environment and need instantaneous access to network resources.
- Satellite offices or branches, ships in the ocean, or teams in remote field locations that need to be connected to a main office or location.

Types of Wireless Networks

Wireless networks can be divided into three categories based on a their technology:

- Local area networks
- Extended local area networks
- Mobile computing

The primary difference between these categories is the transmission facilities. Wireless LANs and extended LANs use transmitters and receivers owned by the company in which the network operates. Mobiles computing use public carriers such as AT&T, MCI, Sprint and the local telephone companies and their public services, to transmit and receive signals.

Wireless communication with Local Area Networks

A typical wireless network looks and acts almost like a cabled network except for the media. A wireless network adapter card with a transceiver is installed into each computer, and users communicate with the network just as if they were at cabled computers.

Access Points

The transceiver sometimes called an access point broadcasts and receives signals to and from the surrounding computers and passes data back and forth between the wireless computers and the cabled network.

These wireless LANs use small wall-mounted transceivers to connect to the wired network. The transceivers establish radio contact with portable networked devices. This is not true wireless LAN because it uses a wall mounted transceiver to connect to standard cabled LAN.

Transmission Techniques

Wireless LANs use five techniques for transmitting data:

1. Infrared
2. Laser
3. Narrow-band (single Frequency) radio
4. Spread-spectrum radio
5. Microwave

Infrared Transmission

All infrared wireless networks operate by using an infrared light beam to carry the data between devices. These systems need to generate very strong signals because weak transmission signals are susceptible to light from sources such as windows.

This method can transmit signals at high rates because of infrared light's high bandwidth. An infrared network can normally broadcast at 10 Mbps.

Four varieties of infrared communications are as follows.

- *Broadband optical telepoint.* This method uses broadband technology. Data transfer rates in this high-end option are competitive with those for a cable-based network.
- *Line-of-sight infrared.* As the name implies, this version of infrared transmits only if the transmitter and receiver have a clear line of sight between them.
- *Scatter infrared.* This technology broadcasts transmissions so they bounce off walls, ceilings and eventually hit the receiver. This has an effective area limited to about 100 feet and has a slow signal because of the entire signal bouncing.
- *Reflective infrared.* In this version of infrared networks, optical transceivers situated near the computers transmit toward a common location, which redirects the transmission to the appropriate computer.

Laser Transmission

High powered laser transmitters can transmit data for several thousand yards when line-of-sight communication is possible. Lasers can be used in many of the same situations as microwave links, but do not require an FCC license. On a LAN scale, laser light technology is similar to infrared technology. Laser light technology is employed in both LAN and WAN transmissions, though it is more commonly used in WAN transmissions.

Narrow-Band Radio Transmission

In narrow-band radio communications (also called single-frequency radio), transmissions occur at a single radio frequency. The range of narrow-band radio is greater than that of infrared, effectively enabling mobile computing over a limited area.

Neither the receiver nor the transmitter must be placed along a direct line of sight; the signal can bounce off walls, buildings and even the atmosphere but heavy walls, such as steel or concrete enclosures can block the signal.

Spread-Spectrum Radio Transmission

Spread-Spectrum radio transmission is a technique originally developed by the military to solve several communication problems. Spread-spectrum improves reliability, reduces sensitivity to interference and jamming and is less vulnerable to eavesdropping than single frequency radio. Spread-spectrum radio transmissions are commonly used for WAN transmissions that connect multiple LANs or network segments together.

As its name suggests, spread-spectrum transmission uses multiple frequencies to transmit messages. Two techniques employed are frequency hopping and direct sequence modulation. Frequency hopping typically transmits at up to 250 kbps, although some versions can reach as high as 2 Mbps. Direct sequence modulation systems operating at 900 MHz support bandwidths of 26 Mbps. Spread-spectrum radio transmissions are often used to connect multiple LAN segments together, thus it is often a WAN connection.

Microwave

Microwave technology has applications in all three of the wireless networking scenarios: LAN, extended LAN and mobile networking. Microwave communication can take two forms: terrestrial (ground) links and satellite links. The frequencies and technologies employed by these two forms are similar, but distinct differences exist between them.

Mobile computing is a growing technology that provides almost unlimited range for traveling computers by using satellite links. The frequencies and technologies employed by these two forms are similar, but distinct differences exist between them.

Three forms of mobile computing are as follows:

- *Packet-radio networking* The mobile device sends and receives network-style packets via satellite. Packets contain a source and destination address and only the destination device can receive and read the packet.
- *Cellular networking* The mobile device sends and receives cellular digital packet data (CDPD) using cellular phone technology and the cellular phone network. Cellular networking provides very fast communications.

- *Satellite station networking* Satellite mobile networking stations use satellite microwave technology.

7.4 Comparisons of Different Wireless Media

The table below compares the different types of Wireless communication media in term of cost, ease of installation, distance and “other issues”

Cable Type	Cost	Installation	Distance	Other Issues
Infrared	Cheapest of all Wireless	Fairly easy, may require line of sights	Under a kilometer	Can attenuate due to fog and rain
Laser	Similar to Infrared	Requires line of sight	Can span several KM	Attenuate due to fog and rain
Narrow-band Radio	Expensive	Requires trained technologies & involves tall radio towers	Span 100 KM	Low-power devices can attenuate & due to fog, rain & solar flares
Microwave	Very Expensive	Involves satellite dishes	Span 1000 KM	Attenuate due to fog, rain and solar flares, can be eaves dropped upon.

7.5 Short Summary

- ☞ Wireless point-to-point communications are another facet of wireless LAN technology. Point-to-point wireless technology specifically facilitates communications between a pair of devices.
- ☞ A wireless bridge acts as a network bridge, merging two local LANs over a wireless connection.
- ☞ Mobiles computing use public carriers such as AT&T, MCI, Sprint and the local telephone companies and their public services, to transmit and receive signals.
- ☞ Telephone cable is UTP and UTP has the highest sensitivity to EMI

- ☞ All varieties of twisted-pair cable have attenuation characteristics that limit the length of cable runs to a few hundred meters, although a 100 meter limit is most common.
- ☞ The fiber optic cables carry data and information and transmit them in the form of light rays.
- ☞ A coaxial cable is capable of transmitting data and information from various workstations simultaneously.

7.6 Brain Storm

1. What are the two types of twisted pair media?
2. What are the names of two common types of coaxial cable and explain briefly?
3. What is a major benefit of fiber-optic cable?
4. What is a major drawback of fiber-optic cable?
5. What is need of wireless network?
6. What are the types of wireless network?
7. What are the some reasons a wireless media would be chosen over a bound media?



Lecture 8

Network Adapter Card

Objectives

In this Lecture you will learn the following:

- Understand the role of the network adapter card in a network including preparing, sending and controlling data
- Knowing the configurable options for network adapter cards.

Coverage Plan

Lecture 8

8.1 Snapshot

8.2 Working of a Network Adapter Card

8.3 Network Adapter Card Compatibility

8.4 Configuring Network Adapter Cards

8.5 Short summary

8.1 Snap Shot

Network adapter cards act as the physical interface or connection between the computer and the network cable. The cards are installed in an expansion slot in each computer and server on the network. After the card has been installed, the network cable is attached to the card's port to make the actual physical connection between the computer and the rest of the network.




Figure 8.1 *Sample Network Adapter Card*

The role of the network adapter card is to:

- Prepare data from the computer for the network cable
- Send the data to another computer.
- Control the flow of data between the computer and the cabling system.

The network adapter card also receives incoming data from the cable and translates it into bytes. Stated at a more technical level, the network adapter card contains the hardware and firmware (software routines stored in read only memory) programming that implements the Logical Link Control and Media Access Control function (in the Data Link layer function of the OSI model).

8.2 Working of a Network Adapter Card

A network adapter card links a PC with the network cabling system. The network adapter card fits into one of the PC's expansion slots. The card has one or more user accessible ports to which the network cabling medium is connected.

Network adapter cards play an important role on the network. They are responsible for translating data from a device on the network mostly computers and converting this data into some form of signal that can be transmitted across the transmission medium. To

enable you to understand the full functionality of the network adapter card, the specific functions of what the network card does must each be addressed.

Preparing and sending data

Before sending network adapter card, it actually sends data over the network, it carries on an electronic dialog with the receiving card so that both cards agree on the following:

- ☞ Maximum size of the groups of data to be sent
- ☞ The amount of data to be sent before confirmation
- ☞ The time intervals between sending data chunks
- ☞ The amount of time to wait before confirmation is sent
- ☞ How much data each card can hold before it overflows
- ☞ The speed of the data transmission

All network cards perform the function of preparing and sending data from a computer to the transmission medium. This data when inside the computer, travels along the bus of a computer in parallel form. This data can move at 8, 16 or 32 bits at a time. The network card must convert these signals coming to it in parallel form, into a serial signal that can travel across the transmission medium. Likewise, when data is received, this serial form of data must be converted into a parallel form matching the bus type (8,16, or 32 bit) being used by the receiving device.

The mechanism of this data conversion is handled in two ways. First, when data is coming from the computer, to be prepared to be sent out on the network, the network adapter card's driver, or software interface, is responsible for converting this data into a format that can be understood by the network adapter card. The standards was NDIS or ODI depending on whether you were going to interface with a Microsoft operating system or a Novell operating system (ODI).

The second part of the data conversion is performed by the physical network card itself. It is here that the actual data that has been passed along from the computer is converted into a serial format using either a digital, analog, or light signal. The network card not only converts the data into this signal but it also responsible for accessing the transmission medium and forming a channel to conduct the signals onto the network. In essence, a network card is like the doorway to the network for the PC or other device.

How a Network Card works

To enable you to fully appreciate how a network card functions, two important concepts must be explained. These are signals and clocking.

Signals

Two basic types of signals are used with transmission media analog and digital.

Analog signals

Analog signals constantly vary in one or more values and these changes in values can be used to represent data. Analog waveforms frequently take the form of sine waves.

The two characteristics that define an analog waveform are as follows;

- Frequency. Indicated the rate at which the waveform changes. Frequency is associated with the wave length of the waveform, which is a measure of the distance between two similar peaks on adjacent waves. Frequency generally is measured in Hertz frequency is illustrated.

- Amplitude. Measures the strength of the waveform.

Each of these characteristics frequency and amplitude can be used to encode data.

Digital Signals

Digital signals are different than analog in that digital signals have two discrete states. These states are either “off” or “on”.

Clocking

Clocking is the mechanism used to count and pace the number of signals being sent and received. Signals are expected to be sent in a continuous flow, representing the start and ending of the data. Clocking is the mechanism used by the network adapter card to determine how much data has been sent. For example. If a network card is designed to transmit data at 20,000 Megahertz a second, others cards receiving this data will also read the data at 20,000 MHz a second. Clocking is a mechanism used by all network adapter cards to measure how much data has been sent or received.

A good example of clocking is when a person taps his feet to keep the time to music. The person doing the tapping expects a set number of music beats per measure; computer network cards also expect so many signals per second.

A clocking mechanism used by some network cards is oversampling. With oversampling, the receiving network adapter card samples, or reads the signals at a higher frequency than that at which the data is sent. This capability is programmed into the card by the manufacturers because the clock used on the sending adapter card can drift apart from that of the receiving adapter card. Oversampling enables the clocking mechanism to determine when this drifting enables the clocking mechanism to determine when this drifting apart is happening so that it can correct the clocking rates.

Measurement of the Signal

To this point you are aware that a network card transmits data and that this data is transmitted between devices across some transmission medium. The network adapter card's role is to convert data from one PC to signals, or convert signals back to understandable data for the PC. These signals are either analog or digital. You also know that clocking is used to count the signals. The last step to understand is the mechanism used by the network adapter card to read the signals. It should be no big surprise at this point that the mechanisms can be grouped into two common methods: digital and analog.

Measurement of Digital Signals

Digital signals use one of two common measurement mechanisms: current state or state transition. The manufacturer builds these measurement capabilities into the network adapter card.

Current State

Current state is a mechanism that uses the clock count to analyze the current state of the signal during that count. This the signal is either "on" or "off" during the clock count.

Changes in the voltages happen during changes in the count. Common digital signal schemes for this sampling mechanism are also known as polar, unipolar, and Biphase.

State Transition

State transition is a more common form of data measurement of digital signal. This form of measurement is used on Ethernet network utilizing copper cables. This form of measurement tends to be less prone to signal disruptions and also does not rely as much on the strength of the signal.

State transition relies on the change of the state of a network signal to represent a new transmission of data. Recall that in current state the length of time a signal is on or off indicates whether the signal represents a 1 or a 0. State transition represents a 1, for example, every time the state of the signal changes on a count, but a 0 is represented every time the state of the signal does not change during a count.

Common state transmission measurement standards are Manchester, Different Manchester, and Biphase Space.

Measurement of Analog Signals

Analog, is like digital signals also follow a similar mechanism of measurement of signals. The main difference between digital and analog signals is that digital have two discrete states “on” and “ off” and analog signals can change frequencies.

Current State

Two mechanism using current state measurement technologies are the frequency shift keying and Amplitude shift keying. FSK uses a change in frequency to indicate a change in data, whereas ASK uses a changes in amplitude to indicate a change in data.

State Transition

State transition of a frequency is the measurement of a frequency’s phase during a clock count. A phase is a difference in transition of a frequency. The transition of a frequency is the change between two frequencies.

An example of phase measurement is that a 1 may be represented by a 90 degree phase shift, and a 0 by no phase shift.

8.3 Network Adapter Card Compatibility

To ensure compatibility between the computer and the network, the network adapter card must:

- ☞ Fit with the computers internal structure (data bus architecture).
- ☞ Have the right type of cable connector for the cabling.

Data Bus Architecture

In the personal computer environment, there are four types of computers bus architectures: ISA, EISA Micro Channel and PCI. Each type of bus is physically different from the others. It is essential that the network adapter card and the bus match.

☞ **ISA (Industry Standard Architecture)**

ISA is the architecture used in the IBM PC, XT and AT computers and all of their clones. It allows various adapters to be added to the system by means of inserting plug-in cards in expansion slots. ISA was expanded from an 8 bit path to a 16 bit path in 1984 when IBM introduced the IBM PC/AT. ISA refers; to the expansion slot itself (an 8 bit slot or a 16 bit). The 8 bit slots are shorter than the 16 bit slots which actually consist of two slots, one behind the other. An 8 bit card could fit into a 16 bit slot, but 16 bit card could not fit into an 8 bit slot.

ISA was the standard personal computer architecture until Compaq and several other companies developed the EISA bus.

☞ **EISA (Extended Industry Architecture)**

This is the bus standard introduced in 1988 by a consortium of nine computer industry companies: AST Research, Inc., Compaq, Epson, Hewlett-Packard, NEC, Olivetti, Tandy, Wyse Technology, and Zenith.

EISA offers a 32 bit data path and maintains compatibility with ISA while providing for additional features introduced by IBM in its Micro Channel Architecture bus.


☞ **Micro Channel Architecture**

IBM introduced this standard in 1988 as part of its PS/2 roll out. Micro channel architecture is electrically and physically incompatible with the ISA bus. Unlike the

ISA bus, the micro channel functions as either a 16 bit or a 32 bit bus and can be driven independently by multiple bus master processors.

☞ PCI (Peripheral Component Interconnect)

This 32 bit local bus used in most Pentium computers and in the Apple Power Macintosh. The current PCI bus architecture meets most of the requirements for providing Plug and Play functionality. Plug and Play both a design philosophy and a set of personal computer architecture specifications. The goal



of Plug and Play is to enable changes to a personal computer configuration with no intervention by the user. The installation of any device should be a simple, fail safe operation. A Microsoft window 96 is a Plug and Play compliant.

Figure 8.2 *ISA, EISA, Micro Channel and PCI Network Adapter Card*

8.4 Configuring Network Adapter Cards

You must configure your operating system so that it can communicate with the network adapter card. Most plug and play adapter cards configure themselves and, in conjunction with the operating system running, assign resources to themselves. In many cases, though you must manually configure the adapter card. These settings are configured through jumper or DIP switch settings, or by using some form of software, so that the network card can communicate with the operating system.

To communicate, the operating system and the network adapter must agree on certain important parameters called resource settings. Some common resource settings for a network adapter are as follows.

- ◆ IRQ
- ◆ Base I/O port address
- ◆ Base memory address

- ◆ DMA channel
- ◆ Boot PROM
- ◆ MAC address
- ◆ Ring speed (token-ring cards)
- ◆ Connector type

IRQ

The IRQ setting reserves an interrupt request line for the adapter to use when contacting the CPU. Devices make requests to the CPU using a signal called an interrupt. Each device must send interrupts on a different interrupt request line. Interrupt request lines are part of the system hardware. The IRQ setting (such as IRQ3, IRQ5, or IRQ15) defines which interrupt request line device is to use. By convention, certain IRQ settings are reserved for specific devices. Different adapter cards provide numerous available IRQs from which you can choose. Be careful, though; some network cards are very limited in the different IRQs that they have available and may offer only IRQs that are currently being used by other devices on your system.

Base I/O port address

The base I/O port address defines a memory address through which data flows to and from the adapter. The base I/O port address functions more like a port defining a channel between the adapter and the processor.

Base memory address

The base memory address is a place in the computer's memory that marks the beginning of a buffer area reserved for the network adapter. Not all network adapter cards use the computer's RAM, and therefore not all adapters require a base memory address setting.

DMA channel

The DMA or Direct Memory channel is an address used for quicker access to the CPU by the adapter card. Many devices, including network cards, often enable you to choose between DMA channels 1 through 7, or have the channel disabled. Not all network cards have the capability to set DMA channels.

Boot PROM

Some network adapter cards are equipped with a Boot PROM. (Programmable Read Only Memory). This Boot PROM enables the network card to boot up and connects over the network, because the Boot PROM has the necessary connection software to use. This feature is often used by diskless workstations, because they have either no hard drives or floppy drives onto which to store connection software.

MAC Address

Using the MAC address burnt into each network card is one of the several ways to establish addresses for nodes on the network. These addresses are hexadecimal in nature and are unique for each card. The IEEE is responsible for assigning these addresses to each network card manufacturer. In some cases, you can reassign a new MAC address for the network adapter card. In the case of ARCNET cards, DIPswitches are used to set the network card's address.

Ring Speed

In the case of token ring networks, the ring speed must be set on the token ring card. The possible values for this are either 4Mbps or 16 Mbps. It is very important that the correct ring speed is set, because an incorrect ring speed prevents your computer from connecting onto the network, or it also can cause the entire network to fail.

Connector Type

Some network cards have different connectors from which you can choose. A common example is an Ethernet card with both a BNC connector and an RJ45 connector. Some network cards require that the connector to be used must be specified. Other network cards self adjust to the connector being used.

Any effort to configure a network adapter card should begin with the card's vendor documentation. The documentation tells you which resource settings are available for you to set, and it might recommend values for some or all of the settings. Some network cards have a default setting, in which all settable values are set to defaults recommended, by the factory.

The actual process of configuring the operating system to interact with a network adapter card depends on the operating system. A plug and play operating system such

as Windows 95, when used with a plug and play compatible Network Adapter Card may perform much of the configuring automatically.

8.5 Short Summary

- ☞ Network adapter cards are the interfaces between the computer and the network cable.
- ☞ To prepare data for the network, the card uses a transceiver to reformat data from parallel to serial transmission.
- ☞ An adapter card has its own unique address.
- ☞ In the personal computer environment, there are four types of computers bus architectures.
- ☞ ISA is the architecture used in the IBM PC, XT and AT computers and all of their clones.

8.6 Brain storm

1. What is the role of adapter card?
2. What type of signal has a discrete state?
3. What are the configurable options on a network card?
4. How can you send and control the data?
5. What is caused by an incorrect setting on a network card?
6. Explain data bus architecture.

☞☞☞

Lecture 9

Connectivity Devices and Transfer Mechanisms

Objectives

In this Lecture you will learn the following:

- Know about Repeaters & Bridges
- About Functions of Repeaters
- Knowing about Hubs

Coverage Plan

Lecture 9
9.1 Snapshot
9.2 Addressing
9.3 Modems
9.4 Repeaters
9.5 Hubs
9.6 Bridges
9.7 Routing
9.8 Gateways
9.9 Short Summary
9.10 Brain Storm

9.1 Snap Shot

An inter network consists of multiple independent networks that are connected and can share remote resources. These logically separate but physically connected networks can be dissimilar in physical type and topology. The device that connects the independent networks together may need a degree of “intelligence” because it may need to determine when packets will stay on the local network or when they will be forwarded to a remote network.

This lecture examines some important connectivity devices. In order to facilitate your understanding of what these devices actually do, the lecture begins by explaining the concept of addressing. This is an important concept to understand in order to differentiate between different connectivity devices that exist on the network. In the following sections, you learn about modems, repeaters, bridges, routers, brouters and gateways. The main emphasis was, where the different connectivity devices were situated within the OSI model. Here the emphasis is more on the function of the different connectivity devices and when one would use them.

9.2 Addressing

Before a discussion of devices is warranted, addressing of network must be explored further. Addressing a network is important, because it is by this mechanism that devices on the network are located and identified.

A network is similar to a city. A city has buildings and it has streets. Now imagine a city in which not two streets had the same name. Each street name must be unique. Now imagine a city that not only had unique name for each street, but also a unique address for each building on its street.

A network is like a street. All buildings on a single street share the same street name that they reside on. Or in other words, all devices on a network segment share the same network address. Thus there are two distinct forms of addresses here: the logical network address and the node address that is physically part of the device.

All packets that go on to a network have within them the source address information and destination address information. Some packets are not routable; therefore they will not contain a source or destination network address in them.

These concepts have been touched on in previous chapters, but it is important that you understand these two concepts here, as the function of different devices on the network is dependent upon device (or hardware) addresses and logical network addresses. As seen in the previous chapter, device addresses in terms of a network card were either burnt onto, or programmed into, the network card by the manufacturer. These addresses could also be set with software by some network card manufacturers, or could be set through DIP switches, as in the case of ARCNet cards. This chapter will address different routing protocols used to discover the existence of different logical addresses on the network.

9.3 Modems

Standard telephone lines can transmit only analog signals. Computers, however store and transmit data digitally. Modems can transmit digital computer signals over telephone lines by converting them to analog form.

Converting one signal form to another (digital to analog in this case) is called modulation. Recovering the original signal is called demodulation. The word "modem" derives from the terms modulation/demodulation.

Modems can be used to connect computer devices or entire networks that are at distant locations. (before digital telephone lines existed, modems were about the only way to link distant devices).

Some modems operate constantly over dedicated phone lines. Others use standard public switched telephone network (PSTN) dial-up lines and make a connection only when one is required.

Modems enable networks to exchange email and to perform limited data transfers, but the connectivity made possible is extremely limited due to the limited bandwidth most modems offer. Modems don't enable networks to connect to remote networks, like a router, to directly exchange data. Instead modems act like network cards in that they provide an access point onto the transmission medium, in this case the telephone lines, in order to send analog signals to another device, most likely another modem, on the network.

Until recently, modem manufacturers used a parameter called baud rate to gauge modem performance. The baud rate is the oscillation speed of the sound wave transmitted or received by the modem. Although baud rate is still an important parameter, recent advances in compression technology have made it less meaningful. Some modems now provide a data transfer rate (in bits per second a more meaningful

measure of network performance) that exceeds the baud rate. In other words, you can no longer assume the baud rate and the data transfer rate are equal.

Modems are classified according to the transmission method they use for sending and receiving data. The two basic types of modems are as follows:

- Asynchronous modems
- Synchronous modems

9.4 Repeaters

As signal travel along a cable, they degrade and become distorted in a process that is called attenuation. If a cable is long enough, attenuation will finally make a sign unrecognizable. A repeater enables signals to travel farther.

Functions of repeaters

A Repeater works at the OSI Physical layer to regenerate the network's signals and resends them out on the other segments.

A repeater takes a weak signal from one segment, regenerates it, and passes it to the next segment. To pass data through the repeater in a usable fashion from one segment to the next, the packets and the Logical Link Control (LLC) protocols must be the same on each segment. This means that a repeater will not enable communication, for example, between 802.3 LAN (Ethernet) and an 802.5 LAN (Token Ring).

Repeaters do not translate or filter anything. For a repeater to work, both segments that the repeater joins must have the same access method. The two most common access methods are CSMA/CD and token passing. That is, they cannot translate an Ethernet packet into a Token Ring Packet.

Repeaters can move packets from one physical media to another. They can take an Ethernet packet coming from a thin net coaxial segment and pass it on to a fiber-optic segment if the repeater is capable of accepting the physical connections.

9.5 Hubs

One network component that is becoming standard equipment in more and more networks is the hub. A hub is the central component in a star topology.




Figure 9.1 A hub is the central point in a star topology

Active Hubs

Most hubs are active in that they regenerate and retransmit the signals, the same way a repeater does. In fact, because hubs usually have eight to twelve ports for network computers to connect, they are sometimes called multiparty repeaters. Active hubs require electrical power to run.

Passive Hubs

Some types of hubs are passive, for example wiring panels or punch down blocks. They act as connection points and do not amplify or regenerate the signal passes through the hub. Passive hubs do not require electrical power to run.

Hybrid Hubs

Advanced hubs that will accommodate several different types of cables are called hybrid hubs. A hub-based network can be expanded by connecting more than one hub.

Hub Considerations

Hubs are versatile and offer several advantages over systems that do not use hubs.

In the standard linear bus topology, a break in the cable will take the network down. With hubs, however, a break in any of the cables attached to the hubs affects only that segment. The rest of the network keeps functioning.

Other benefits of hub-based topologies include:

- Changing or expanding wiring systems as needed. Simply plug in another computer or another hub.
- Using different ports to accommodate a variety of cabling types.
- Centralized monitoring of network activity and traffic. Many active hubs contain diagnostic capabilities to indicate whether or not a connection is working.

9.6 Bridges

Like a repeater, a bridge can join segments or workgroup LANs. However a bridge can also divide a network to isolate traffic or problems.

For example if the volume of traffic from one or two computers or a single department is flooding the network computers or that department.

Bridges can be used to:

- Expand the distance of a segment
- Provide for an increased number of computers on the network
- Reduce traffic bottlenecks resulting from an excessive number of attached computers.

A bridge can take an overloaded network and split it into two separate networks, reducing the amount of traffic on each segment and making each network more efficient.

- Link unlike physical media such as twisted pair and coaxial Ethernet
- Link unlike network segments such as Ethernet and Token Ring, and forward packets between them.

Functions of Bridges

Bridges work at the Data Link layer of the OSI model. Because they work at this layer, all information contained in the higher levels of the OSI model is unavailable therefore, they do not distinguish between one protocols which pass across bridges, it is up to the individual computers to determine which protocols they can recognize.

You may remember that the Data Link layer has two sub layers, the Logical Link Control sub layer and the Media Access Control sub layer. Bridges work at the Media Access Control sub layer and are sometimes referred to as Media Access Control Layer Bridge.

A Media Access Control layer bridge:

- ☞ Listens to all traffic.
- ☞ Checks the source and destination addresses of each packet.
- ☞ Builds a routing table as information becomes available
- ☞ Forwards packets in the following manner.

If the destination is not listed in the routing table, the bridge forwards the packets to all segments, or if the destination is listed in the routing table, the bridge forwards the packets to that segment (unless it is the same segment as the source).

A bridge works on the principle that each network node has its own address. A bridge forwards packets based on the address of the destination node.

Bridges actually have some degree of intelligence in that they learn where to forward data. As traffic passes through the bridge, information about the computer addresses is stored in the bridge's RAM. The bridge uses this RAM to build a routing table based on source address.

Initially, the bridge's routing table is empty. As nodes transmit packets, the source address is copied to the routing table. With this address information, the bridge learns which computers are on which segment of the network.

Bridge Considerations

Bridges have all of the features of a repeater, but also accommodate more nodes. They provide better network performance than a repeater. Because the network has been divided, there will be fewer computers competing for available resources on each segment.

To look at it another way, if a large Ethernet were divided into two segments connected by a bridge, each new network would carry fewer packets, have fewer collisions, and operate more efficiently. Although each of the networks was separate, the bridge would pass appropriate traffic between them.

Remote Bridges

Bridges are powerful tools in expanding and segmenting a network, they are often used in large network, that have widely dispersed segments joined by telephone lines.

Only one bridge is necessary to link two cable segments. However, in a situation where two separate LANs are located at a great distances from each other, they may be joined into a single network. This can be done by implementing two remote bridges connected with synchronous modems to a dedicated, data grade telephone line.




Figure 9.2 Remote bridges can be used to connect remote segments

Because remote LAN segments can be joined over telephone lines, there may be a situation where multiple LANs are joined by more than one path. In this situation, it is possible that data might get into a continuous loop. To handle this possibility, the IEEE 802.1 Network Management Committee has implemented the spanning tree algorithm (STA). Under STA, software can sense the existence of more than one route, determine which would be the most efficient, and then configure the bridge to use that one. Other paths are disconnected using software, though the disconnected routes can be reactivated if the primary route becomes unavailable.

Difference between Bridges and Repeaters

Bridges work at a higher OSI layer than repeaters. This means that bridges have more intelligence than repeaters and can take more data features into account.

Bridges are like repeaters in that they can regenerate data, but bridges regenerate data at the packet level. This means that bridges can send packets over long distances using a variety of long distance media.

9.7 Routing

An internetwork consists of two or more physically connected independent networks that are able to communicate. The networks that make up an internetwork can be of very different types. For example, an internetwork can include Ethernet and token ring networks.

Because each network in an internetwork is assigned an address, each network can be considered logically separate that is, each network functions independently of other networks on the internetwork. Internetwork connectivity devices, such as routers, can use network address information to assist in the efficient delivery of messages. Delivering packets according to logical network address information is called routing. The common feature that unites internetwork connectivity devices (routers and brouters) is that these devices can perform routing. The following list details some common internetwork connectivity devices:

- Routers
- Brouters

Each of these devices is discussed in the following sections.

Routers

In an environment consisting of several network segments with differing protocols and architectures, a bridge may not be adequate for ensuring fast communication among all of the segments. A complex network needs a device, which not only knows the address of each segment, but can also determine the best path for sending data and filtering broadcast traffic to the local segment. Such a device is called a router.

Routers work at the Network layer of the OSI model. This means they can switch and route packets across multiple networks. They do this by exchanging protocol specific information between separate networks. Routers read complex network addressing information in the packet and, because they function at a higher layer in the OSI model than bridges, they have access to additional information.

Routers can provide the following functions of a bridge:

- ☞ Filtering and isolating traffic
- ☞ Connecting network segments.

Routers have access to more information in packets than bridges, and use this information to improve packet deliveries. Routers are used in complex network situations because they provide better traffic management than bridges and do not broadcast traffic. Routers can share status and routing information with one another and use this information to bypass slow or malfunctioning connections.

You can use routers to divide large, busy LANs into smaller segments, much as you can use bridges. But that's not the only reason to select a router. Routers also can connect different network types. An example of this would be a router that connected a token-ring segment with the Ethernet segments. On such networks, a router is the device of choice as a bridge cannot perform this function.

The protocols used to send data through a router must be specifically designed to support routing functions. IP, IPX and DDP (the AppleTalk Network-layer protocol) are routable transport protocols. NetBEUI is a non-routable protocol transport protocol.

Because routers can determine route efficiencies, they usually are employed to connect a LAN to a wide area network (WAN). WANs frequently are designed with multiple paths and routers can ensure that the various paths are used most efficiently.

Routers come in two general types :

- **Static:**

Static routers do not determine paths. These routers require an administrator to manually set up and configure the routing table and to specify each route.

- **Dynamic:**

Dynamic routers do an automatic discovery of routes and therefore have minimum amount of set up and configuration. They are more sophisticated in that they examine information from other routers and make packet by packet decisions about how to send data across the network.

Difference between static and dynamic routers

Static Routers	Dynamic Routers
Static routers are considered more secure because the administrator specifies each route.	Security can be improved in dynamic routers by manually configuring the router to filter out network addresses discovered and prevent traffic from going there
Manual set up and configuration of all routes.	Manual configuration of the first route. Automatic discovery of additional networks and routes.
Always uses the same route which is determined by a routing table entry	Can choose a route based on factors such as cost and amount of link traffic
The route used is hard coded and is not necessarily	Can decide to send packets over alternate routes.

Brouters

A brouter is a router that also can act as a bridge. A brouter attempts to deliver packets based on network protocol information, but if a particular Network layer protocol isn't supported, the brouter bridges the packet using device addresses.

9.8 Gateways

The term "gateways" originally was used in the Internet protocol suite to refer to a router. Today the term "gateway" more commonly refers to a system functioning at the top levels of the OSI model that enable communication between dissimilar protocol systems. Gateways make communication possible between different architectures and environments. They repackage and convert data going from one environment to another so that each environment can understand the other environment data. A gateway repackages information to match the requirements of the destination system. Gateways can change the format of a message so that it will conform to the application program at the receiving end of the transfer. For example, electronic mail gateways, such as the X.400 gateway, receive messages in a one format, translate, it and forward in X.400 format used by the receiver, and vice versa. Gateways commonly function at the OSI Application layer, but actually can operate at any level of the OSI model.

A gateway links two systems that do not use the same:

- ☞ Communication protocols
- ☞ Data formatting structures
- ☞ Languages
- ☞ Architecture

Gateways interconnect heterogeneous networks, for example, Microsoft Windows NT Server to SNA (IBM's systems Network Architecture). They change the format of the data to make it conform to the application program at the receiving end.

Gateways can be implemented as software, hardware or a combination of both. An example of a gateway is often seen in email systems. When you send email, say from Microsoft Exchange to some on the Internet, a gateway is responsible for converting the Microsoft Exchange message contents and addressing, to one that is compatible with the SMTP (Internet) message format and addressing.

9.9 Short Summary

- ☞ Repeaters operate at the physical layer of the OSI model. Repeaters do not rely on any type of address at all, as their function is to simply regenerate a signal, regardless of where the signal is destined to go.
- ☞ Bridges operate at the Data Link Layer of the OSI model. Bridges analyze the hardware addresses or MAC addresses, of the source and destination device on each packet. Based upon these MAC addresses, a packet will be allowed through a bridge if that device is on the other side, or not through a bridge if the device is not on the other side of the bridge.
- ☞ Routers utilize logical network addresses to define separate network segments. These logical network addresses are usually assigned by the network administrator and will follow specific naming conventions, depending upon the protocol being used.
- ☞ Routers, brouters and bridges are designed to isolate traffic not regenerate signals to exceed cable distance recommendations.
- ☞ Hubs operate at the physical layer of the OSI model.

9.10 Brain Storm

1. Define Repeaters, Bridges
2. When discussing Repeaters, Bridges and Routers, which device relies on logical network addresses, physical hardware addresses, or does not use addresses at all?
3. What are the functions of Repeaters?
4. Describe the differences between Repeaters and Bridges?
5. Describe the difference between static and dynamic routing
6. Write the advantages of hubs & explain the types of hubs?
7. You need to connect a token-ring and an Ethernet LAN segment. To do so you will need what type of device?



Lecture 10

File Systems

Objectives

In this Lecture you will learn the following:

- Knowing about File Systems
- About type of File Systems
- Understanding NTFS permission

Coverage Plan

Lecture 10

10.1 Snapshot

10.2 File System

10.3 Understanding NTFS Permissions

10.4 Compressing NTFS Files and Folders

10.5 Short Summary

10.6 Brain Storm

10.1 Snap Shot

Administrating security include implementing, planning and enforcing policy for protecting data and shared network resources, which include directories, files, printers, plotters. In this lecture you will learn about File Systems & types of file system.

10.2 File System

A file system provides persistent storage of information. The file system interfaces with disk drives and provides a system of organizing the way information is stored in the tracks and clusters of the drives. User interface with the file system by working with files and directories , both of which can have various attributes such as read-only, read/write, execute and so on.

A local file system may allow users to only access local disk drives. However, most operating systems also allow users to access disk drives located on other network computers and to share or “publish” directories on their own systems. These network-aware operating systems require a higher level of security since unknown users may access the files over the network. Therefore, more advanced network operating systems like Windows NT and NetWare provide special attributes to access account and must be properly authenticated before thy can access a file.

Most file systems today use the hierarchical directory structure in which a directory tree is built from a root directory. Directories are like containers that can hold other directories (subdirectories) or files. A directory is a parent to its subdirectories. Directories have attributes that are usually “ inherited” by all the files and subdirectories of the directory; however attributes can be changed of individual files and directories in most cases

Types of File System

The file systems are of various types. A network operating system usually features more than one file system. Windows NT supports a number of OS. Windows NT supports four types of file systems. They are

- FAT
- VFAT
- NTFS
- CDFS

FAT

The FAT file system is the only file system supported by Windows NT where floppy disk drives are used. FAT file system was the default file system on personal computers. The reason for this was that the DOS was the dominant operating system used the FAT for file management on these computers and the great evolution and revolution has occurred.

Contiguous	Sector	Storage locations on media that are in consecutive locations
Cluster		The basic unit of space allocated on a disk. A cluster may contain one or more sectors
Directory		A logical, hierarchical form of file organization used on various forms of media

Fragmented File This refers to a file that is stored on a disk in non contiguous sectors

The following factors are important in implementing a FAT file system for Windows NT

- ❁ FAT has minimal file system overhead(less than 1MB)
- ❁ FAT is the most efficient file system for partitions smaller than 400 MB
- ❁ Performance declines with large numbers of files because the FAT file system uses a linked list for the folder structure. Also if the amount of data in a file grows, the file can become fragmented on the hard disk, which makes the process of retrieving the file from the disk slower
- ❁ Windows NT prevents a deleted file from being undeleted. Undelete utilities require direct access to the hardware, and Windows NT does not allow user mode applications to directly access hardware. However, if the deleted file is on a FAT partition and the system is restarted under MS-DOS it may be possible to undelete the file if it has not been written over.

FAT Naming Conventions

Under Windows NT, the FAT file system is enhanced to support long file names(LFNs). The following criteria apply to file names on a Windows NT partition that has been formatted with the FAT file systems:

- The name can be up to 255 characters, including the full path
- The name must start with either a letter or a number and can contain any characters except the following: “/ \ [] : ; | = , ^ * ?”
- The name can contain multiple spaces
- The name can contain multiple periods. The characters after the last period are treated as the extension.
- Names preserve case but are not case-sensitive.

Security

A FAT partition cannot be protected by the local file or directory security features of Windows NT. The only security available on FAT partitions is provided through Windows NT directory-level sharing mechanism.

VFAT File System

Virtual File Allocation Table (VFAT) is an extension to FAT that allows newer programs and operating systems to use long file names. The long file names can be up to 255 characters. They can contain spaces as well as periods, which are not supported by FAT. VFAT file names are always stored with a FAT alias.

NTFS File System

NTFS file system is the native file system for Windows NT. Unlike FAT, NTFS is not limited to a fixed number of sectors per cluster. In the NTFS file system the cluster is the base unit. The cluster factor is defined as a number of bytes and formatting a volume as NTFS guarantees that the cluster factor is a multiple of the sector size on the drive. As NTFS guarantees that the cluster factor is a multiple of the sector size on the drive. As NTFS address everything by a cluster number, the file system is isolated from consideration of the sector size of the underlying drive. Due to this isolation from the physical characteristics of the disk, the number of sectors per cluster is a default suggested value rather than a hard and fixed value. NTFS allows modification of the default suggested value rather than a hard and fixed value. NTFS allows modification of the default number of sectors per cluster to better suit the actual usage of the

volume. NTFS also seeks contiguous disk space before writing or copying a file to the disk. Because of this, simply copying an NTFS volume off to another disk and back actually helps defragment the disk. This anti-fragmentation feature of NTFS drives work best when they are at less than 60 percent capacity. Over that amount NTFS is forced to fragment its files.

The Windows NT File System (NTFS) is the file system supported only by Windows NT. Both, windows NT Server and Windows NT Workstation support NTFS. No other Windows or non-Windows Operating System, supports NTFS file system

When you have to install Windows NT for providing very high security, format the partition using NTFS. This is always the case when you are installing a Windows NT server in a domain. You can also prefer NTFS format to secure your Windows NT Workstation from local unauthorized access.

New Technology File System (NTFS) is the proprietary file system of Windows NT. NTFS is much more secured than the FAT, and hence is the preferred choice over FAT> NTFS is a better over FAT because of the following features.

NTFS Naming Conventions

The following rules apply to NTFS file names:

- ❁ File and folder names can be up to 255 characters, including extensions
- ❁ In general, names are not case-sensitive but are case-preserving. There is an exception: a file name is case-sensitive when it is used by a POSIX application. NTFS allows the coexistence of two identically named files that differ only in case
- ❁ Names can contain characters except the following: ? “ / \ < > * | :

Fault Tolerance:

NTFS allows system logs to be recorded in a log file, which helps in troubleshooting. It detects bad sectors and moves the data stored on these bad sectors to another good area on the disk. It then marks those bad clusters so that the data will not stored on them in future.

Security:

NTFS provides a very high degree of security. FAT file system does not provide only directory permissions but the NTFS security provides file permissions in addition to the directory permissions.

NTFS provides security for file server and clients in a corporate environment. NTFS supports access control and ownership privileges important for the integrity of corporate data. A user logged on locally to a Windows NT server featuring FAT file system can access all files and directories on that server irrespective of the permission assigned to him. However, a user logged on to a windows NT server featuring NTFS file system, can access only those files and directory to which he is given they access permissions.

Additional features

NTFS has additional features that make it a powerful and flexible file system. Other NTFS features include:

- ◆ **Performance** - NTFS decrease fragmentation, as it sees that the files are saved in contiguous blocks. This increases the performance largely.
- ◆ **Compression** - NTFS provides a great degree of compression. A folder can be compressed with all its subfolders, which makes a great amount of disk free space on the server hard disk. File compression reduces text-oriented application files by about 50 percent and executable files by about 40 percent
- ◆ **Transaction – based recoverability** NTFS has high reliability. It is a recoverable file system that used transaction logging to log all folder and file updates automatically. This is used by Windows NT to redo or undo operations that failed due to system failure, power loss and so on.
- ◆ **Bad-cluster remapping** - NTFS supports the recovery technique of bad-cluster remapping. If an error occurs because of a corrupt sector on the hard disk, NTFS allocates a new cluster to replace the cluster with the corrupt sector. NTFS then marks the original sector as corrupt. this is transparent to any applications performing disk I/O
- ◆ **Support for Macintosh files** - Installing services for Macintosh o computers running Windows NT server allows Macintosh computer files to be stored on the computer running Services for Macintosh

- ◆ **Support for posix requirements** - NTFS is the POSIX.1 complaint supported file system and supports the following POSIX.1 requirements:
 - **Case-sensitive naming** - Under POSIX file names are case-sensitive, so that README.TXT, Readme.txt and readme.txt are different files
 - **Additional time stamp** - This supplies the time that the file was last accessed

There may be situations in which another file system in addition to NTFS is necessary.

NTFS Partition Boot Sector

The partition boot sector is located at the beginning of the NTFS volume, with a duplicate located elsewhere on the volume. Starting at sector 0, the boot sector may be as large as 16 sectors. The contents of the boot sector are

- ☞ The BIOS parameter Block (BPB)
- ☞ The number of sectors in the volume
- ☞ The starting Logical Cluster Numbers(LCNS) of the Master File Table(MFT) and the Master file Table Mirror (MFT2)

A duplicate of the boot sector is located at the end of the volume on Windows NT 4.0 and up, and at the logical center of the volume for Windows NT 3.51 and earlier versions.

The NTFS system or metadata files are created on the volume during the format of the volume by the FORMAT command. The NTFS system files are all defined with fixed file numbers in the MFT. These metadata files, such as the Bad cluster file, which tracks the location of all identified bad clusters in the partition, follow immediately after the MFT in the root directory of the partition. After the last metadata file are the data files. The MFT2 is not a complete mirror of the MFT, but mirrors the first three records of the MFT (the MFT,MFT2 and the log file) to prevent a bad cluster developing at the beginning of the MFT from corrupting the partitions.

NTFS Implementation Consideration

The following considerations are important when implementing NTFS

- Recoverability is designed into NTFS. Users do not have to run a disk repair utility on an NTFS partition.
- NTFS provides security for files and directories but does not provide file encryption.
- NTFS maintains a separate Recycle Bin for each user, which increases security. If a user deletes a confidential file, no other user can log on and pull the confidential file out of the Recycle Bin
- When using removable media, NTFS does not support changing hard disks without restarting the computer
- It is not possible to format a floppy disk with NTFS because of the overhead involved in NTFS.
- Fragmentation is greatly reduced on NTFS partitions
- If it is necessary to defragment a file the file can be copied to another drive, and then copied back to the original drive. When the file is copied back to the original drive, NTFS attempts to place it in a contiguous block.
- If a file increases in size after it is on the drive, it may potentially become fragmented, depending on the drive's space usage.

Converting a FAT Partition to NTFS

It is possible to convert an existing FAT hard disk partition to NTFS

Windows NT includes the convert command which can be used to convert FAT partitions to NTFS. Note that conversion is a one-way process; there is no way to convert an NTFS volume to a FAT partition. Converting a partition from FAT to NTFS preserves all data on the partition, unlike formatting the partition, which destroys all data.

The convert command uses the following command line syntax: convert drive: F (Here F is the partition of Hard disk)

CDFS (CD-ROM File System)

This file system is native to NT. It supports what has traditionally been called long file names. Long file name supports enable CDs to be read only, hence making them good candidates for boot media.

10.3 Understanding NTFS Permissions

Share-level access control provides only a limited capability to determine which users can access which files. The FAT file system offers only share-level access control. If you need to control access down to subfolders and individual files, your only choice is to use the NTFS file system. Doing so is no sacrifice at all, because NTFS offers more features, better performance (on all but the smallest volumes), and better security.

In addition to the file name, file size, and data/time stamp, NTFS stores extended attributes with each file and folder entry. One extended attribute, *permissions*, determines which users and groups have access to the shared resource. NTFS has the following types of permissions.

- *File-access permissions* store information about which users and groups are permitted access a specified file and the level of access they're allowed. For example, the user Admin and the group Programmers might have full read/write access to a particular database file; the group Marketing might have read only access; and the group Accounting might have no access at all.
- *Folder-access permission* store information about which users and groups are permitted access a specified file and the level of access they're allowed. For example, the user Webmaster and the group Administrators can have full read/write access to the Web server folder on your server, the group Everyone can have read-only access, and the user Guest can have no access at all.

By default, users inherit file and folder permissions from the group of which they are members. For example, if a newly created user is assigned to the group *marketing*, he is automatically granted all the group's file and folder -access permissions. If a user is a member of more than one group, he has *all* permissions owned by *any* group of which he is a member.

Paying careful attention to how you assign group file and folder permissions allows you to reduce or eliminate the time-consuming and error-prone process of assigning permissions on a user-by-user basis.

Working with NTFS File-Access Permissions. NTFS file-access permissions control which users and groups can access a file, and at what level. Remember that NTFS file-access permissions can further restrict the access level granted by share permissions, but they can't extend access beyond that granted by share-access permissions. You can add, modify, view, or remove the following file-access permissions for each file:

- ❖ *No Access (None)* permission restricts all access to the shared file.

- ❖ *Read (RX)* permission allows you to view the file name and open the file in read-only mode, but you can't write to the file or delete it. Because read (R) permission implies execute (X) permission, if the file is an executable program file, read permission allows you to execute it.

- ❖ *Change (RWXD)* permission grants all the rights provided by Read permission and adds the rights to write (W), execute (X) and delete (D) the file, create new files and subfolders, modify the contents of new or existing files, and delete files and subfolders.

- ❖ *Full Control (All)* permission grants all the rights provided by Change permission and adds the rights to change NTFS file-access and folder permissions, as well as take ownership of NTFS files and folders.

- ❖ *Special Access* permission allows you to customize the file-access permissions for a particular file. You can specify any combination of read (R), write (W), execute (X), delete (D), change permissions (P), and take ownership (O). For example, you can use Special Access file-access permissions to allow a specified user or group to have read, write, and execute permissions for the file, but not to have delete permission.

Modifying, Viewing, and Removing NTFS File-Access Permissions. Follow these steps to modify, view, and remove NTFS file-access permissions:

1. In Windows NT Explorer, select the file(s) for which permissions are to be added, modified, viewed, or removed.

2. Right-click to display the context-sensitive menu, and choose Properties to display the *Filename* Properties sheet.
3. Click the Security tab to display the Security page.
4. Click the Permissions button to display the File Permissions dialog.
5. Select a type of access from the Type of Access drop-down list. You can choose one of the standard types of access - No Access, Read, Change, or Full Control - or you can select Special Access to customize file-access permissions for this file or group of files.
6. If you've selected one of the standard access types, click OK to apply the selected file-access permissions. You then return to the *Filename* Properties sheet. Click OK again to accept the changes and close the *Filename* Properties sheet.

If you select the Special Access, Mark the check boxes to select the types of access to be granted for the selected file(s).

7. After you select the permissions for the file, click OK to accept these settings and return to the File Permission dialog.
8. In the File Permission dialog, click OK to apply the selected file-access permissions and return to the *Filename* Properties sheet. Click the OK again to accept the changes and exit the *Filename* Properties sheet.

Adding NTFS File-Access Permissions Follow these steps to add NTFS file-access Permissions:

- 1 Follow steps 1 through 4 from the preceding section to display the File permissions dialog.
- 2 Click the Add button to display the Add Users and Groups dialog.
- 3 Select the domain or computer from which the users and groups are to be added from the list names from drop-down list. The Names list displays available groups . To display individual users within these groups, click the show users button.

- 4 Select individual users or groups that you want to add file-access permissions for help double-clicking the name in the Names list. Each appears in the Add Names list as you select it.
- 5 Select the access type to be granted to the selected users and groups from the Type Access drop-down list.
- 6 Click OK to accept your changes and return to the File Permission dialog. The new added users or groups appear in the names list. If you need to assign Special File Access permissions to the newly added users or groups do so now by using the steps description in the preceding section and then return to the File Permissions dialog.
- 7 Click OK to return to the filename Properties sheet , Click OK again to accept the changes and exit the filename properties sheet.

Working with NTFS Folder-Access Permissions

NTFS folder-access permissions control which users and groups can access a folder and its files, and at what level. Remember that NTFS folder access permissions can further restrict the access level granted by share-access permissions.

You can add, modify, view or remove the following access permissions for each folder. Each named permission affects the folder in question and the files contained within it. The first parenthetical item after each access permission name lists the effect of that permission on the folder; the second parenthetical item lists the effect of that permission on files contained within the folder.

- No Access (None) permission restricts all access to the shared folder, specifying No Access for a user eliminates his access to the folder, even if he is a member of a group or groups that have access to the folder.
- List (RX) permission allows users to view a list of files and subfolders contained within the folder, and to change to a subfolder, but it doesn't grant permission to access the files.

- Read (RX) permission grants all the rights provided by List permission. It allows users to open a file in read-only mode, but not to write to the file or delete it. Because read (R) permission implies execute (X) permission, if the file is an executable program file, read permission allows you to execute it.
- Add permission allows users to create new files and new subfolders within the folder, but doesn't grant permission described in the preceding paragraphs.
- Add & Read permission combines the rights granted by the Read and Add folder permission described in the preceding paragraphs.
- Change (RWXD) permission grants all rights provided by the Add & Read permission and adds the rights to write and delete files and subfolders.
- Full Control permission grants all the right provided by the change permission and adds the rights to change NTFS file access and folder permission, as well as take ownership of NTFS files and folders.
- Special Directory Access permission lets you customize folder access permissions. You can specify any combination of read (R), write (w) execute (x) delete (d) change permissions (p) and take ownership (o). for example, you can use Special Directory Access folder access permissions to allow a specified user or group to have list and read permissions for files within the folder, but not to have execute permission.
- Special File Access permission allows you to customize file access permission. You can specify any combination of read (R) write (W) execute (x) delete (D) change permission (P) and take ownership (O) . Special file access permission works in the same way as the Special Directory Access permissions but affects only specified files contained within the folder rather than the folder itself.

Modifying, Viewing, and Removing NTFS folder Access Permissions.

You can modify, view, and remove NTFS folder access permissions by following these steps.

1. In Windows NT Explorer, select the folder(s) for which permissions are to be added, modified, viewed or removed.
2. Right-click to display the context sensitive menu, and choose Properties to display the folder name properties sheet.
3. Click the Security tab to display the Security page.
4. Click the Permission button to display the Directory Permission dialog.
5. Select an access type from the Type of Access drop down list. You can choose one of the standard types of access No Access, List, Read, Add, Add & Read, Change or Full Control. You also can choose special directory access to specify a custom set of access rights for the affected folders, or Special File Access to specify a custom set of access rights for the files contained within those folders.
6. If you have selected one of the standard access types, click OK to apply the selected folder access permission. You then return to the folder name properties sheet. Click OK again to accept the change and exit the Foldername properties sheet.

If you select Special Directory Access, the Special Directory Access dialog opens. Mark the check boxes to select the types of access to be granted for the selected folder or folders.

7. After you select the permission for the folder, click OK to accept these settings and return to the Directory Permissions dialog

8. In the Directory Permission dialog, click OK to apply the selected folder access permissions and return to the Foldername Properties sheet. Click OK again to accept the changes and exit the Folder name Properties sheet.
9. If you select Special File Access, the Special File access dialog opens. Mark the check boxes to select the types of access to be granted for files contained within the selected folder or folders.
10. After you select special file access permissions for the affected folder or folders click to accept these settings and return to the Directory Permissions dialog.
11. In the Directory Permissions dialog click OK to apply the permissions and return to the Foldername properties sheet. Click OK again to accept the changes and exit Foldername Properties sheet.

Adding NTFS folder-Access Permissions

Follow these steps to add NTFS folder-access permissions:

1. Follow steps 1 through 4 from the preceding section to display the Directory Permissions dialog.
2. Click the Add button to display the Add Users and Groups dialog.
3. Select the domain or computer from which the users and groups are to be added from the List Names from drop-down list. The Names list displays available groups. Click the show Users button to display individual users from within these groups.
4. Double-click the names in the names list to select individual users or groups for which you want to add file-access permissions. Each appears in the Add Names list as you select it. (You can also select multiple users and groups in the Names list by using standard Windows conventions for making multiple selections) After you finish making selections, click the Add button to transfer all selected names to the Add Names list.

5. Select the access type to be granted to the selected users and groups from the Type of Access drop-down list.
6. Click OK to accept your changes and return to the Directory Permissions dialog. The newly added users or groups are displayed in the Name list. If you need to assign special Directory Access permissions or special File Access permissions, follow the steps in the preceding section.
7. After you properly assign all permissions, return to the Directory Permissions dialog and click OK to return to the Foldername Properties sheet. Click OK to accept the changes and exit the Foldername Properties sheet.

10.4 Compressing NTFS Files and Folders

You can save a substantial amount of disk space by compressing files shared from folders on NTFS partitions. Text, work processing, spreadsheet, and some static graphics files usually achieve better than 40 percent compression. NTFS's file compression/decompression system is similar to that of MS-DOS 6.0's DoubleSpace and MS-DOS 6.22's DriveSpace but is more reliable than either product. The major benefit of NTFS compression is that you can selectively compress individual folders or files; DoubleSpace and DriveSpace require that you compress an entire disk volume.

Using Windows NT Explorer's or My Computer's Compression Features

It's most common to use Windows NT Explorer or My Computer to compress entire shared NTFS folders or individual files. To compress a folder and all its files with Explorer follow these steps:

1. Launch Explorer and choose Options from the View menu. Click the View tab and mark the Display Compressed Files and Folders with Alternate Color Check box.
2. Select the folder you want to compress.

3. Choose Properties from the File menu to display the General page of the Foldername properties sheet.
4. Mark the Compress check box and click OK to open the Explorer dialog.
5. Mark the Also Compress Subfolders check box. The Explorer dialog opens whether or not the folder you're compressing has subfolders.
6. Click OK to continue the compression process. The Compress File message box displays the compression status. As mentioned earlier in the section "Compressing NTFS Files and Folders" compression ratios of 40 percent (a 100k file compresses to 60K) are common for text, word processing, worksheet, and similar files.

On completion of the compression process, Explorer displays in blue the file name, size, date, and other text entries for the compressed folder and its files.

To Compress an entire volume, follow this procedure:

1. Select the drive letter in Explorer and choose Properties from the File menu to open the volume's Properties sheet.
2. Mark the compress check box of the General page and Click OK.
3. Mark the Also Compress Subfolders check box and click OK.

Moving and Copying Compressed Files

Compressed files, including folders carry a compressed (C) attribute. Following are the general rules that govern whether a compressed file is compressed or decompressed when moved or copied:

- When you move a compressed file from one NTFS folder to another NTFS folder (including moves between NTFS volumes), the file remains compressed whether or not the destination drive or folder is compressed.

- When you copy a compressed file to a NTFS folder, the compression status of the file corresponds to the compression attribute of the folder. That is , copying a compression file to a decompressed folder results in a decompressed copy; copying a decompressed file to a compressed folder results in a compressed copy.

- When you move or copy a compressed file to a folder of a FAT volume, the resulting file is compressed.

Using Compact.exe at the Command Prompt

Compact.exe is a command-line utility that you can execute in a batch file to compress or decompress volumes, folders and files. Following is the command-line syntax for Compact.exe obtained by redirecting compact /? To a text file:

```
COMPACT [/C|/U] [/S[:dir]] [/A] [/I] [/F] [/Q] [filename [...]]
```

/C Compresses the specified files. Directories will be marked so those files added afterward will be compressed

/U Uncompress the specified files. Directories will be marked so that files added afterward will not be compressed

/S Performs the specified operation on files in the given directory and all subdirectories. Default "dir" is the current directory.

/A Displays files with the hidden or system attributes, these files are omitted by default.

/I Continues performing the specified operation even after errors have occurred. By default, COMPACT stops when an error is encountered.

/F Forces the compress operation on all specified files, even those which are already compressed. Already-compressed files are skipped by default.

/Q Report only the most essential information. Filename specifies a pattern, file, or directory.

Multiple command-line parameters must be separated by spaces. If you omit the parameter, Compact.exe displays the compression state of the current folder and the files contained in the folder. Unlike Explorer and My computer, compact.exe doesn't ask whether you want to compact subfolders; it automatically compacts all subfolders and their files.

10.5 Short Summary

- ☞ Directories are the folders, which has multiple files. The files of the same categories are grouped together.
- ☞ NTFS provides a very high degree of security.
- ☞ FAT32 an enhanced version of the FAT file system available on some version of Windows 95 is not supported by Windows NT
- ☞ Share level permissions cannot give you more permission than granted by the NTFS permissions
- ☞ FAT is the only file system supported by Windows NT where floppy disk drives are used

10.6 Brain Storm

1. Explain Files and Directories in NT file system.
2. Explain the file systems that are all supported by NT file structure.
3. What are the four types of file system, explain briefly?
4. Discuss Share level protection in NT.
5. NTFS file system can access Fat and CDFS – True or False Explain?
6. “If NT server uses Fat file system then all the files and directories will have NTFS Rights”- Discuss?.



Lecture 11

Computer Security

Objectives

In this Lecture you will learn the following:

- Understanding the concept of computer security
- Standards of Computer security- C2
- Knowing about the requirements for C2 security

Coverage Plan

Lecture 11

- 11.1 Snap Shot
- 11.2 Security and Windows NT
- 11.3 C2 Security
- 11.4 Requirement for C2 security
- 11.5 Short summary
- 11.6 Brain Storm

11.1 Snap Shot - Computer Security

Broadly defined, computer security involves prevention of undesired or unintentional access to any part of a computer system. In an inclusive sense, this would include all aspects of hardware and software, servers work stations, LAN devices and cabling as well as computer operating systems, network operating systems, user programs, and the most important element user data.

11.2 Security and Windows NT

It's hard to have a coffee table conversation about Windows NT without at least broaching the subject of security. Security is built into the core of Windows NT and represents an integral part of Microsoft marketing strategy for their NT product line.

UNIX has been around for over two decades, and the number of books written about security from a UNIX point of view is over whelming. Windows NT, on the other hand, has only been around a few years, and there is still little extensive information published that focuses directly on Windows NT specific security issues , such as the way that NT handles its user security database and network authentication. Reading a generalized treatise on computer and network security is useful for developing ideas that can be extended to the NT environment, but there are critical NT specific issues that need to be addressed directly. Although there are architectural similarities between Windows NT and various flavors of UNIX, the complexities of computer security are directly linked with the intricacies of the operating system, such as how users are created, where and how the user database is stored, how you access files and what privileges different processes have.

Windows NT has some wonderfully robust features that allow for very secure installations, however, like any complex system, there are many things that can go wrong. NT system can be compromised if you don't fully understand the security illicitness of your actions or even in some cases your look of actions.

Although many users still consider computer security to be something that they don't need to be concerned with, they couldn't be further from the truth. The increasing appearance of computer networks in today's corporate environment requires that users begin to get more involved with this issue.

With stand-alone computer it is relatively easy to identify you potential attackers an develop means to protect yourself. However as soon as you attach that computer to a network, you've just changed the nature of the game.

11.3 C2 Security

Listed prominently on Windows NT's list of features is that it meets the C2 computer security guidelines set by the U.S government.

The National Computer Security Center (NCSC) published the two definitive texts on computer security. Department of Defense with regard to qualifying computer system for use in various critical environments. Although these guidelines must often be adhered to in many government agencies.

The Department of Defense's Trusted Computer System Evaluation Criteria (TCSE) also called orange book for the color of its cover provides security requirements for automatic data processing systems. The companion to this book called the Red Book Extends the interpretations for the evaluation of trusted systems to compute networks. These books were developed to provide users a metric for determining the degree of trust they can put in their compute system.

Users are held accountable for their actions. System tracks all process and records actions on a user-by-user basis. Prevents object reuse, and most ensure efficacy of system security monitor. Users can grant others access to their data.

11.4 Requirements for C2 Security

We've identifies that Windows NT is in fact a product that has been evaluated to meet the C2 security rating based on the NCSC Orange Book. Let's review on more time the key criteria for C2 class systems

- **Discretionary control** The operating system must enable the owner of an object (in most cases a file, print job, or process) to grant or restrict access to the resource. This discretionary control feature is implemented in Windows NT through access control lists, which provide single user granularity for controlling object access.
- Object reuse discarded or delete objected must not be accessible, either accidentally or purposefully. Because all object assignments must pass through a single point, the Security Reference Monitor the system guarantees that discarded objects cannot be accessed by any other process. For example when a file is deleted under the NTFS, the file's data cannot be retrieved, hence the lack of an undelete facility. Additionally, when a process is allocated memory, Windows NT initializes this space, ensuring that no data from a prior process will be left behind.

- **User identification and authentication** Before being allowed to access the system, each user must identify him or herself by entering a unique logon name and password. The system uses this data to validate a user's right to access the system. If the credentials are valid, the user is granted a security token, which is used to validate any object access during that particular logon session. Furthermore, the user's logon information can be used to track actions performed by the users.
- **Auditing** An integrated feature of Windows NT is its capability to audit all object access and securities related actions on the system and identify which user performed the actions. This capability to provide such a granular level of auditing is necessary for C2 compliance. Additionally, access to this audit data is restricted to authorized administrative users.
- **Protection** Windows NT prevents processes from accessing memory outside its own 32 bit space, there by removing the possibility that an application can read or write data from another process's memory space. The kernel itself runs in a protect 32 bit memory space. Processes can communicate with the kernel and with other processes through the use of a carefully defined message passing system, which removes any need for a process to modify another process's memory. Additionally, because Windows NT prevents direct hardware access, n process can bypass the security model and directly access memory or hard disk information

11.5 Short Summary

- ☞ Computer Security involves prevention of undesired or unintentional access to any part of a computer system
- ☞ Red book – A book that provides security requirements for automatic data processing systems
- ☞ Divisions of the Trusted Computer System Evaluation Criteria

Class	Protection Level
D	Minimal Protection
C1	Discretionary Protection
C2	Controlled Access Protection
B1	Labeled Security Protection
B2	Structured Protection
B3	Security Domains
A1	Verified Design

11.6 Brain Storm

1. When we need Computer Security?
2. What is C2 Security?
3. What are the benefits of C2 security?

☞☞☞

Lecture 12

Windows NT Server Installation

Objectives

In this Lecture you will learn the following:

- About Installation and Compatibility issues
- Requirements for installing Win NT

Coverage Plan

Lecture 12

- 12.1 Snap Shot – Installing Window NT server
- 12.2 System Requirements
- 12.3 Compatibility Issues
- 12.4 Types of Installation
- 12.5 Starting the Basic Installation
- 12.6 Repairing the Windows NT Server Operating System Installation
- 12.7 Short Summary
- 12.8 Brain Storm

12.1 Snapshot - Installing Windows NT Server

The installation program is an application that will do the work of installing the network operating system in a variety of ways, depending on:

- The environment into which it will be installed
- The size of the network
- The types of jobs the server will perform in the network
- The type of file system the server will use.
- The identification of the server.
- Operating systems in the server.
- How the server's hard disk space is divided

The installation program asks a series of question to determine the installation parameters. The following text describes some of the key question areas.

12.2 System Requirements

Understanding the system requirements of Windows NT server 4.0 is the first step in the planning process. The Following summarizes the requirements across the supported hardware platforms.

Category	Intel Platforms
Processor type	Intel 486/25 minimum; Intel Pentium recommended
Number of CPUs	4 maximum out of the box; Support for more CPUs available From hardware vendors
Memory	12 MB minimum, 16 MB recommended
Hard Disk	130 MB minimum Space
Floppy Disk Drive	3.5-inch (5.25-inch no longer supported)
CD-ROM Drive	Required if not installing over a network
Monitor	VGA or higher resolution
Pointing Device	Mouse recommended
Network Adapter	Required for network connectivity Card
Optional Components	Modems / ISDN cards for remote Connectivity

12.3 Compatibility Issues

The hardware platforms supported by Windows NT Server 4.0 come in various makes and models. It is very important that you review the Windows NT Hardware Compatibility List (HCL) included with your Windows NT Server 4.0 package to ensure that your particular make and model is supported. The HCL is also available as a Windows help file (HCL.HLP) in the Support directory of the Windows NT Server 4.0 CD-ROM. For the latest additions to the HCL, visit Microsoft's Web site at HTTP:

[//www.microsoft.com/ntserver/hcl/hclintro.htm](http://www.microsoft.com/ntserver/hcl/hclintro.htm). If you plan to use any of the following, confirm that the hardware is supported in the HCL:

- SCSI and other controller cards
- Video and other graphics cards
- Tape and other backup devices
- Network cards
- CD-ROM drives
- UPS
- PCMCIA cards

Due to the proliferation of new computer products, the HCL may not contain information about one or more of your particular computer hardware or software components. In this case, you should check with the manufacturer of your product to determine compatibility with Windows NT Server 4.0

After you have determined that your computer hardware is supported by Windows NT Server 4.0, the next step is to get the latest software driver for the product.

There is a high probability that your computer vendor has posted new drivers and information on the Internet since the release of the product to Windows NT Server 4.0. Usually, the updated information provides tips to work around known problems, and the new drivers will contain enhancements that fix problems that you may experience with Windows NT Server 4.0

12.4 Types of Installation

For new installation, Windows NT supports different installation types for Intel platforms. The preferred method for installing Windows NT for the first time on RISC platforms is from the CD-ROM. The following table summarizes the options.

Installation	Intel Platforms (Install Files in the I386 Directory)
Floppy-based	<p>Boot your computer with the Windows NT Server4.0 startup disk # 1 and follow the instructions.</p> <p>OR</p> <p>Run WINNT-EXE from the CD-ROM to automatically create start-up floppy disks and continue installation.</p>
CD-ROM only	<p>Connect to your CD-ROM drive and run WINNT.EXE B. This option copies temporary files to your hard disk and begins installation.</p>
Network Distribution	<p>Connect to an MS-DOS based network path and Run WINNT.EXE. This command may be combined with other switches for floppy less installation and/or automated installation.</p>
Automated	<p>Connect to your distribution files using any of the preceding methods and run the following WINNT.EXE /u:<script>, where <script> is the script file that contains all the settings that you want to use for your installation.</p>

12.5 Starting the Basic Installation

After you have connected to your installation files and typed the appropriate command to begin installation, Windows NT Server 4.0 proceeds in typically the same fashion (except as noted in the previous Table) for all installation methods across all supported platforms. The steps are as follows.

1. Setup detects mass storage devices and allows you to add other devices not detected.
2. Setup displays the Windows NT Server 4.0 license screen and other legal information that you must accept to continue installation.
3. Setup displays the hardware components found and allows you to accept or change the information
4. Setup asks you to select a partition and a directory in which to install Windows NT Server 4.0
5. Setup begins to copy files from your distribution point.

6. Setup finished the DOS portion of the installation and restarts the computer to continue the Graphical User Interface (GUI) portion of the installation. Windows NT Server 4.0 is started (in installation mode) at this point
7. Setup copies more files and begins to gather additional information about your computer
8. Setup prompts for personalization of your installation of Windows NT Server 4.0
9. Setup prompts you to set user license information.
10. Setup prompts you to select a unique name for your Windows NT Server 4.0 computer.
11. Setup asks you to determine the role that this installation of Windows NT Server 4.0 will play on your network.
12. Setup prompts you to assign a password to the default administrator account.
13. Setup asks if you wish to create an emergency repair disk.
14. Setup prompts you to add or remove software components.
15. Setup initializes for networking and prompts you to select network or Remote Access Services (RAS) adapter.
16. Setup asks if you wish to install the included Internet Information Server (IIS).
17. Setup asks you to select a network or other connectivity adapter cards
18. Setup prompts you to select network protocols.
19. Setup allows you to install network components other than the standard selection
20. Setup installs components
21. If selected, setup prompts you to configure TCP/IP
22. Setup displays binding information so that you may modify it if necessary.
23. Windows NT starts the network
24. Windows NT prompts you to select, create, or join a domain
25. Setup asks you to set the clock and time zone.
26. Setup lets you test your video adapt and add a newer driver if necessary.
27. Setup finalizes, configures your computer to run Windows NT, and copies programs and apples.
28. Windows NT sets security on system files.
29. Windows NT saves your configuration information and asks to create the emergency disk (ERD) if you have chosen this option.
30. Windows NT installation completes and your computer restarts.

Running the Setup Program

If you're installing windows NT server 4.0 on a computer with a CD-ROM drive supported by windows NT server, insert the first Setup disk, labeled Setup Boot Disk, in the A; drive. Then insert the windows NT server CD-ROM in the CD-ROM drive, reboot the computer, and skip to the section "Choosing to install or Repair". Setup uses disks for first part of the installation procedure, and then automatically detects and uses the CD-ROM drive for the rest of the installation.

The following sections pertain to installing windows NT Server over the network from another server or from a CD-ROM drive that is not supported by Windows NT server but is supported by DOS or Windows 95.

Creating setup Boot Disks. The Windows NT Server 4.0 distribution CD-ROM claims to be bootable, but relatively few CD-ROM drives or system BIOSes support bootable CD-ROMs. In most cases, you receive a Can't Find NTLDR message if you try to boot from the CD-ROM. To install Windows NT server 4.0 on most PCs with a supported CD-ROM but no operating system installed, you must use Setup boot disks to install a minimal version of Windows NT on the PC. If you have the distribution CD-ROM but no boot disks, have three formatted, blank 3 ½ inch disks ready and follow these steps to create the required boot disks;

1. Insert the distribution CD-ROM in the CD-ROM drive of a PC running DOS, Windows NT, or windows 95. if the setup splash screen appears, close it.
2. At the command prompt, log on to the CD-ROM drive and change to the 1386 folder (or the folder for the appropriate platform).
3. Type Winnt /ox (DOS or Windows 95) or Winnt 32/ ox (windows NT) and press Enter to start the Setup program.
4. Accept or correct the path to the CD-ROM files, and then press enter or click continue. Windows NT display the dialog.
5. Insert a disk labeled windows NT 4.0 Server Setup Disk #3 in the A: drive and click OK to continue.
6. When prompted, insert the remaining two disks, Windows NT 4.0 Server Setup disk #2 and Windows NT. 4.0 Server setup Boot Disk, and press Enter or click OK.

The disks are necessary if you Windows NT installation becomes corrupt, requiring of the emergency repair disk created later in the installation process. Leave the setup boot disk in the A: drive and restart you computer, booting from the A: drive. Skip to the "choosing to install or repair" section to continue with the description of the setup process.

Installing from a network server. Installing over the network is more common for workstations, which may not all have CD-ROM drives, than for servers. You must have an operating system installed and a functional network connection to install Windows NT server from a network server. To install windows NT server 4.0 over the network, you need another machine on the network with either the Windows NT server CD-ROM in a shared CD-ROM drive, or a prepared Windows NT server installation folder with a copy of the three folders from the CD-ROM.

- \1386 contains all the files needed to install to computers that use Intel 1386 family.
- \Mips is for RISC machines using silicon Graphics MIPS processor(s).
- \Alpha is for RISC machines using DEC's Alpha processor family.
- \Ppc is for PowerPCs from Apple, IBM, and others that use the Power PC RISC processor.

To prepare an installation folder for network setup of Windows NT server, follow these steps.

1. Create a folder on the server with an appropriate name, such as installnt.
2. Share the folder, giving the administrator group at least Read access.
3. Create a subfolder with the same name as the one you copy from the CD-ROM for example, \1386.
4. Copy the installation files from the CD-ROM's subfolder to the new installation subfolder.

To start setup from the new computer after the Windows NT installation folder is created, or to run from an unsupported CD-ROM drive, you follow almost identical steps.

1. Label four blank 3 1/2 inch disks as setup boot disk, setup disk 2, setup disk 3, and windows NT Server Emergency Repair Disk. Add the name of the computer to the emergency repair disk. You can use the other three disks for installation on other computers; the emergency repair disk is specific to the computer on which the installation is made.
2. Start the machine where Windows NT server is to be installed by using the existing operating system, which must either support an attached CD-ROM drive or have a network connection.

3. In the machine is using DOS, change to the network drive and directory or to the CD-ROM drive and directory that holds the installation files. If the machine is using windows 95, open an Explorer window for the network or CD drive. If the machine is using a previous version of windows NT, open a File Manager window or CD drive. If the machine is using a previous version of Windows NT, open a File Manger window for the network or CD drive. (if the machine runs windows 3.1x, exit windows if necessary and perform the installation from DOS.)
4. From DOS, run WINNT.EXE by typing winnt and pressing Enter. Form earlier versions of Windows NT, run Winnt32 exe by double clicking the file in File Manager. From Windows 95, run winnt.exe by choosing Run from the Start menu, typing winnt, and then pressing enter.
5. When requested provide the drive and directory name for the location of the installation files.
6. Insert each formatted setup disk when prompted. The WINNT or WINNT 32 program copies the required setup files to the disks.
7. Continue the setup process as described later, starting with the section "Choosing to Install or Repair".

Installing from an Existing Operating System

If you have a compatible CD-ROM drive and operating system, you can start the setup program from DOS or windows as follows.

- For windows NT 3.5 + or DOS/windows 3.1+, use the method described earlier in the "creating Setup Boot Disks" section, but omit the /ox command line parameter.
- For windows 95, inserting the CD-ROM in the drive uses AutoPlay to display the splash screen. (this splash screen also appears if you insert the CD-ROM into a PC running Windows NT 4.0 Which also supports AutoPlay.

In either case, you are prompted to confirm the location of the files on the CD-ROM; then Setup copies all the setup files to a temporary folder, requests you

to reboot, and begins the installation process. You need approximately 250M of free disk space to store the installation copies and the working copies of the files. When the files are copied, remove the CD-ROM from the drive (and the disk, if any, in you're A: drive), and reboot the computer. On restarting, the Boot.ini file automatically starts the Windows NT setup process.

Choosing to Install or Repair

If you use the setup disks, Setup runs from the setup boot disk, requests a second disk, and then displays a blue (DOS) screen titled Windows NT server 4.0 Setup: Welcome to Setup, which asks whether you want to repair an existing Windows NT server installation or continue with a full installation.

The first part of this lecture covers installing Windows NT server 4.0 for the first time, so press Enter to continue the installation process. You are requested to insert Setup Disk #2 and press Enter.

Setting up you Fixed Disk Drives

After you watch the detection of mass storage devices, add more drivers if needed, and then confirm you hardware, you next must make Windows NT server work with you fixed disk drive(s) this involves the following basic steps.

1. Partition the drive(s)
2. format with the right file system.
3. specify the installation folder.

The following sections describe each of the preceding steps in details.

Partitioning the Drive(s). The first partition is the system partition. On RISC machines, the system partition must be drive C; has sufficient room for the system files (about 148M). As noted earlier in this chapter, you can elect to install the non-bootable windows NT system files in another partition, if necessary, to reduce the C: partition disk space requirement.

Setup shows the partitions that already exist on your fixed disk drive. Use the up and down arrow keys to highlight a partition or the unpartitioned space; then press D to delete it, C to create a partition in unpartitioned space, or enter to choose it as the system partition.

Formatting with the Right File System. Next, choose a file system and optionally format the partition with that system. You can choose to format the partition with FAT or NTFS; keep in mind, however, that these two options destroy the data stored on the partition. If the partition is already formatted, you can elect to convert it to NTFS or retain the existing format. As noted earlier, unless you have a compelling reason to use the FAT format, choose NTFS. Use the up and down arrow to highlight your choice (for example, Convert to NTFS) and press Enter to continue.

Specifying the install folder. Setup now needs to know where to put the windows NT and system files. The default suggestion, C:\Winnt, ordinarily is satisfactory. You can, however, install these files in a folder of any name on a partition of the first physical drive. To change the folder name, press backspace to remove the suggested name and type your chosen name.

At this point, setup looks previous versions of windows on your machine. (if you formatted you system partition in the previous step, you wiped out any previous version that were on the machine.) Rather than ask you to choose whether to upgrade or dual boot, setup determines the answer based on whether you use the same folder. Following are your options when you choose the same or a different installation folder.

- If you have a previous versions of windows NT, installing in the same folder causes you old setting to be used in this new installation and the old version to be removed. Installing in a different folder allows you to dual boot between the older version and windows NT server 4.0.
- If you have Windows 95, you can't migrate your setting and aren't allowed to use the same folder. You can dual-boot between Windows 95 and Windows NT server 4.0
- If you have windows 3.x you dual boot regardless of the choice you make. If you use the same folder your settings are migrated; if you use a different folder, you must specify new settings throughout the rest of Setup.

Watching the Copying Process

Now that setup can access your partitioned and formatted drive(s), and has established where the system files are to go, it's time to copy the files to the

fixed disk. Before doing so, setup offers to examine your fixed disk drive for defects. Press enter to allow the examination to proceed, or Esc to skip it. (if you are having trouble installing windows NT server and find yourself at this point in the install repeatedly, it's not necessary to repeat the examination every time). After the examination is performed or skipped, Setup copies files to the folder chosen in the previous step. The time required to create the copies depends on the speed of the CD-ROM or network connection, and the performance of your fixed disk drive.

When the copying process is complete, windows NT server is ready to run, but additional configuration information is required to finalize the installation process. At this point, you are prompted to remove a disk still in the A: drive. Also remove the CD-ROM. Press Enter to restart the system.

After the computer restarts, the rest of the process proceeds under the operating copy on windows NT server. Setup's simple character based interface is replaced with windows NT style dialogs, and you indicate choices by clicking dialog buttons instead of pressing keys. The help and Exit setup options remain available but are activated by dialog command buttons. Also, you use the back and Next buttons to move through the Setup wizard (as Microsoft calls which next stage of Setup). you are requested to insert the CD-ROM into the drive indicated in the copy Files From text box, and then click OK to copy the additional files. After the files are copied, click Next to continue.

Fine-Tuning the Installation

If the computer is to be a dedicated server rather than a workstation, you can save disk space by skipping the installation of accessories such as CD player, games, screen savers wallpaper, and so on.

Chose the components you want installed or not installed in the Select Components dialog. If you want to install part of a component, such as multimedia, click the details button to select individual applications. Click OK on the Details dialog to return to the list of components. When that list has each component selected, deselected, or partially selected as you prefer, click Next to move to the next stage of Setup.

Joining the Network

Now all of setup is complete, except for setting up the network. Confirm that you want to go on to the network portion of setup by clicking next. You can't get back to earlier screens after moving to the network portion. You perform the following general steps for the network installation;

1. Describe your connection to the network
2. Choose to install Internet information server
3. Choose and configure a network adapter.
4. Choose protocols.
5. Choose services
6. Confirm network bindings
7. Start the network and join a domain.

Describing your connection to the network. The first question setup asks is how you connect to your network. If your machine has a network adapter card (as it almost certainly does), choose wired to the network. If you dial up to a network (an unlikely choice for Windows NT Server) choose Remote Access to the network.

Choosing to install Internet information server. Internet information Server makes your information available over the Internet or a corporate intranet. At this point in the setup process, you specify only whether you intend to install it.

Choosing and Configuring a Network Adapter. Assuming that you have one or more network adapters, in the next step setup detects them automatically. A functioning network adapter card is required for installation of a primary domain controller or a backup domain controller, but the network doesn't need to be operational at this point in the installation process.

Click start search to find the first adapter. If a second adapter is to be found, click find next to search for it. If you have an adapter that wasn't found, click select from list to specify the card yourself. Click next to move on.

Choosing protocols. Next, choose one or more network protocols. IPX/SPX, TCP/IP, or NetBEUI. You can configure multiple protocols by checking more than one box in this dialog. The most common combination is NetBEUI and TCP/IP, unless you have an existing Novell NetWare network, in which case

you should select all three protocols. You can add or remove network protocols with control Panel's Network tool after installing windows NT server.

Choosing services. You can choose to install any of these five network service.

- internet information server
- RPC configuration
- NetBIOS interface
- Workstation
- Server

Adding new services after the fact is harder than adding them now, and you can use control panel's services tool later to disable services that you don't want to run. You also can add or remove network services with the services with the services page of Control panels Network tool. For most server installations, it's best to install all the preceding services at this point.

Even more services are available if you click choose from list.

- DHCP relay Agent
- Gateway (and client) services for NetWare
- Microsoft DHCP server
- Microsoft DNS Server
- Microsoft TCP/IP Printing
- Network Monitor Agent
- Remote Access Service
- Remoteboot service
- RIP for internet protocol
- RIP for NWLink IPX/SPX compatible transport
- RPC support for Banyan
- SAP Agent
- Services for Macintosh
- Simple TCP/IP services
- SNMP Service
- Windows internet name services

Network Setting. Setup confirms that you're ready to install the adapters, protocols, and services that were selected over the previous few dialogs. Click

Next, and you have the opportunity to confirm adapter settings, such as the interrupt number and I/O port address. As discussed earlier in the lecture you should know these setting before you start the installation procedure.

The Network settings property sheet is next. You can display the same properties by using the protocols page of control panel's Network tool after you install Windows NT Server 4.0. To configure a protocol, select it from the list on the protocols page and click properties button. If you're not sure how to use this property sheet you can leave it for now and bring it up again after windows NT server is completely installed.

Each network protocol is configured separately. The following figure shows the configurations dialog TCP/IP the Microsoft TCP/IP Properties sheet.

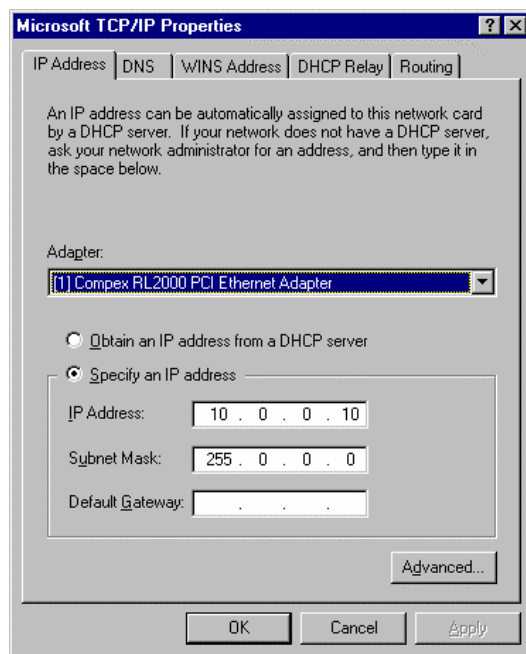


Figure - The IP address page of the Microsoft TCP/IP Properties sheet

On the IP address page, set the IP address and subnet mask for your machine, or tell windows NT to use DHCP (dynamic Host Configuration protocol) to assign an IP address dynamically. If you didn't establish these setting during your network planning process, get them from the person who did.

The DNS (domain name service) page is used to control the way the server looks up domain names of other computers on a LAN or WAN. You enter the IP addresses of one or more DNS servers accessible to your server. DNS

servers translate a fully qualified domain name such as www.mcp.com into an IP address. You can get these addresses from the same person who provided your server's IP address. Some networks use multiple DNS servers, checking the local one first and then asking a remote DNS server if the name wasn't found in the local one. To change the priority of an IP address within the list, click the UP or Down button.

The WINS address page describes the way that wins looks up the domain names of other computers. Enter the IP addresses of your primary and secondary wins servers as provided by the person who told you your own IP address. For machines local to your network, you may want to use DNS and LMHOSTS services; if so select their check boxes. The Scope ID is usually left blank provide a scope ID only if you're told to do so by your network administrator.

The DHCP relay page identifies your DHCP servers. These servers manage IP addresses within your internal network, and you get their addresses from the person who told you to use DHCP rather than specify the IP address of your machine. If your need to adjust the other parameters on this page rather than accept the defaults, this person will tell you so.

The routing page is relevant only to machines with more than one network adapter, each with its own IP address(es). If you turn on IP forwarding, your server can route traffic between the two networks.

Confirming Network Bindings Now setup gives you a chance to adjust your network bindings. If you aren't sure what network bindings are or why you might want to adjust them, leave them alone. The default bindings usually are adequate for installation and network startup.

You can enable and disable communications between services and adapters or protocols by double clicking the service name, clicking the adapter or protocol, and then clicking the enable or disable button.

Starting the network and joining a domain After windows NT is configured for your network hardware. Setup loads the network software and establishes a connection to the network. Confirm that you've made all the choices by clicking next and then wait while the network starts. Setup asks for a domain name. Provide the same domain name you used earlier and wait while setup

searches the network. If you're joining a domain but not installing a primary domain controller, provide the administrator name and password here.

Finishing Server Setup

Click the finish button to move to the final setup steps. Setup created program menu groups and desktop icons, if you had an earlier version of Windows installed. Adding groups and icons doesn't install the applications or make any changes to the settings and configurations stored in windows 3.1 's. INI or REG.DAT files, or in windows 95's registry. You must return the application's setup program if the application must be reconfigured for windows NT or if you are upgrading from Windows 95.

Installing Internet Information Server To install only World Wide Web services, clear the Gopher and FTP service check boxes. Click OK to continue. Accept the default folders for the service(s) you install, unless you have a specific reason to do otherwise. Click OK twice to continue and then to confirm your choices. Setup copies a number of files to your drive. Click OK when advised that you must establish an Internet domain name for the server.

In the Install Drivers dialog, select SQL server in the available ODBC drivers list, even if you haven't installed SQL server, and click OK to continue.

Setting the Time and Date You now get a chance to set the time, date, and time zone. On the Time Zone page, select your time zone from the drop-down list. This list is arranged numerically according to the time difference from Universal Time (UT, formerly known as Greenwich mean time, GMT). Zones east of Greenwich, England, appear above it in the list. After you choose your time zone, the map of the world in this dialog scrolls so that your time zone is in the center. Windows NT Server knows the rules for use of Daylist Savings Time; make sure that the Automatically Adjust for Daylist savings Time box is marked or cleared, as appropriate for your location.

Configuring Your Display. Until now, the setup program has used standard VGA resolution with 16 colors. Windows NT attempts to detect the type of chip on your graphics adapter card. If you are using an adapter with a popular Windows graphics accelerator chip for which a DriectDraw driver is included with Windows NT, click OK when the Detected display message box appears.

Otherwise, you must install a driver from a disk provided by the graphics card supplier.

The settings page of the Display Properties sheet lets you set a resolution and color depth suited to the combination of your graphics adapter card and video display unit. Some adapter cards provide additional features and controls on the settings page. For a simple adapter, follow these steps to set your display properties;

1. Click the Display Type button to confirm that the correct graphics adapter cards has been detected and change the selection, if necessary.
2. Adjust the number of colors, resolution, and refresh frequency. For servers, a resolution of 800*600 pixels and a color depth of 256 colors is adequate. If you have a 15 inch or smaller display, the large fonts selection improves readability in 800*600 resolution.
3. Click the Test button to examine the of the settings you choose.
4. Click OK to close the Display properties sheet.

Creating the Emergency Repair Disk

At this point, windows NT server 4.0 is installed, configured, and ready to act as a network server. The final step is to create an emergency repair disk to use in case of a catastrophic failure. Any 3 1/2 inch disk will suffice, because setup formats the disk before copying the files. Label the emergency repair disk with the server name you assigned and store it in a safe location.

Be sure to update your emergency repair disk frequently to keep the configuration data up to date. Your server's configuration is likely to change appreciably during the first few hours of use as your install applications, change users and groups, and so on. The pace of change slackens with time, but you should make a habit of updating your emergency repair disk regularly. To create an updated emergency repair disk, follow these steps:

1. Insert you original emergency repair disk or a new disk in drive A: from the Start menu choose Run, type `rdisk` in the Open text box, and click OK to open the Repair disk Utility dialog .

2. Click Update Repair info to update the disk's content. (if you have lost your emergency repair disk, insert a blank disk and choose create repair disk.)
3. Confirm your intent to update the existing emergency repair disk. Click OK when asked whether you want to create an emergency repair disk, and then click OK to format the disk and create the updated version.
4. After the updated emergency repair disk is created, click Exit to close the Repair Disk Utility dialog.

It is a generally accepted practice to create a duplicate of the emergency repair disk for off site storage, along with duplicated of driver disks you used during setup. Driver disks also are needed during the repair process.

Restarting the server for the last time.

After you make the emergency repair disk, setup completes all the tasks required to install and configure windows NT server 4.0 and automatically restarts the server. Remove the emergency repair disk from the A: drive and remove the CD-ROM to allow the system to boot from the fixed disk drive.(the boot process takes more time if you converted a FAT partition to NTFS as part of the setup process, because it is at this point that the format conversion occurs). After windows NT loads, press Ctrl+Alt+Delete to log on, and use the administrative account name and password you created earlier to log on to windows NT server. Your installation is complete.

12.6 Repairing the Windows NT Server Operating System Installation.

You can run your windows NT server software in various ways, and in some cases there's no way to recover. Making complete backups frequently is one way to reduce the recovery work involved; having an emergency repair disk is another.

This section assumes that you have a problem that prevents your server from booting successfully. It further assumes that the last known good choice during the boot process is of no assistance, and that you can't edit the Registry

remotely from another computer to adjust your settings. The last resort is to reinstall windows NT server from scratch but first try using your emergency repair disk.

Gather the emergency repair disk, originals or copies of the three setup disks, any driver disks you used during the original installation, and the original CD-ROM used for installation. If any of these are missing, the repair process may be impossible. Follow these steps to attempt to repair the server.

1. Correct hardware problems, if any, and install replacement hardware as needed.
2. Boot from the windows NT server 4.0 boot disk, switch to setup Disk 2 when prompted, and select repair at the first prompt, you can select one or more of these choice;
 - Inspect Registry Files
 - Inspect Startup environment
 - Verify windows NT system files
 - Inspect boot server

Move the highlight up and down through this list with the arrow keys, and press enter to select or deselect the highlighted item. If you have no idea what's wrong, leave all the items selected. When your list is complete, highlight continue and press enter.

3. Repair checks for mass storage devices by using the same process as Setup. The steps and keystrokes involved are identical to those discussed in the earlier section "Detecting Mass storage".
4. When asked whether you have the emergency repair disk, press Enter if you do and Esc if you don't.
5. You're prompted for the original installation media so that installed files can be compared to the originals.
6. After a partial file examination process, you have a chance to restore Registry files. You should try at least one repair attempt without restoring Registry files. If that doesn't make the system bootable, and no backups are

available, repeat the repair process and restore the Registry files. All changes made since the emergency repair disk was created whether you installed an entire application suite or changed a user's desktop settings will be lost.

7. Repair examines the remaining files on your fixed disk drive and compares them to files on the installation media. If a file is found that varies from the original, you're given four choices.
 - Press Esc to skip this file that is, leave the version Repair thinks is corrupted on your fixed disk drive.
 - Press enter to repair this file that is, copy the original file from the CD-ROM to your fixed disk drive.
 - Press A to repair this file and all additional files Repair thinks are corrupted, with no more prompts.
 - Press F3 to stop the repair attempt.
8. When the Repair process is complete, you're prompted to remove any disk still in the drive. Press Enter to restart the computer.

When a reboot is successful, you have the following problems to tackle;

- You must establish the cause of your catastrophic problem and make sure that it has been solved.
- You will likely need to perform some reconfiguration because some of your configuration information might be lost or reverted to old values.
- You must restore data from backup tape(s) if data files were lost.
- You must notify users of the problem and let them know whether they need to update their own configurations.

If you can't reboot the system and suspect corruption of your MBR or partition boot sector, use the techniques outlined in the next section.

If the Repair process doesn't work after several tries, you must reinstall windows NT server from scratch, and then reinstall applications as needed. In the event of serious corruption, you may lose all the user and group

information for your server. A successful Repair operation takes about half as long as an install; always try a Repair first.

12.7 Short Summary

- ☞ Windows NT Server 4.0 provides many installation options. Advance planning is the key to efficient use of the Setup program
- ☞ You can install Windows NT server 4.0 supports from CD-ROM in less than 30 minutes, if you know the pitfalls to avoid.
- ☞ Microsoft recommends at least a 300M FAT system partition for systems that don't require full operating system security. This space is used for Windows NT installation, pagefile and DOS or Windows 95 installation.
- ☞ Windows NT doesn't support Windows 95's Plug and Play(PnP) features of ISA adapters.
- ☞ Using TCP/IP, requires your address to be unique across your LAN or WAN

12.8 Brain Storm

1. What are the System Requirements for Installation?
2. How do you Install Windows NT workstation?
3. What are the various Types of Installation?
4. What is HCL?



Lecture 13

Setting of RAID

Objectives

In this Lecture you will learn the following:

- About RAID Levels
- Knowing to create Windows NT Server Stripe and Mirror sets
- Able to recover the RAID Levels

Coverage Plan

Lecture 13

13.1 Snap Shot

13.2 Understanding RAID Levels

13.3 Creating Windows NT Server Stripe and Mirror sets

13.4 Recovering a Software RAID 1 or RAID 5 Set

13.5 Short Summary

13.6 Brain Storm

13.1 Snap Shot

A redundant array of inexpensive disks (RAID) uses multiple fixed-disk drives, high-speed disk controllers and special software drivers to increase the safety of your data and to improve the performance of your fixed-disk subsystem. All commercial RAID subsystems use the Small Computer System Interface (SCSI, pronounced scuzzy), which now is undergoing a transition from SCSI-2 to SCSI-3 (also called Ultra Wide SCSI). Virtually all network servers now use narrow (8-bit) or wide (16-bit) SCSI-2 drives and controllers. Ultra-wide SCSI host adapters for the PCI bus can deliver up to 40M per second (40M/sec) of data to and from the PC's RAM.

RAID levels 1 and higher protects your data by spreading it over multiple disk drives and then calculating and storing parity information. This redundancy allows any one drive to fail without causing the array itself to lose any data. A failed drive can be replaced and its contents reconstructed from the information on the remaining drives in the array. There are several levels of RAID, each of which is optimized for various types of data handling and storage requirements. You can implement RAID in hardware or in software.

This lecture describes, the different types of RAID levels and how to create the Window NT Server Stripe and Mirror set.

13.2 Understanding Raid Levels

Disk drives do only two things; write data and read data. Depending on the application, the disk subsystem might be called on to do frequent small reads and writes or the drive need to do less frequent, but larger, reads and writes. An application server running a client server database, for example tends toward frequent small reads and writes, whereas a server providing access to stored images tends toward less frequent, but larger, reads and writes. The RAID levels vary in their optimization for small reads, large reads, small writes, and large writes. Although most servers have a mixed disk access pattern, choosing the RAID level optimized for the predominant environment maximizes the performance of your disk subsystem.

The various RAID levels are optimized for various data storage requirement in terms of redundancy levels and performance issues. Different RAID levels store data bit wise, or sector-wise over the array of disks. Similarly parity information might be distributed across the array or contained on a single physical drive.

RAID levels 1 and 5 are very common in PC LAN environment. All hardware and software RAID implementations provide at least these two levels. RAID level 3 is used occasionally in specialized applications and is supported by most hardware and some software RAID implementations RAID levels 2 and 4 are seldom, if ever, used in PC LAN environments, although some hardware RAID implementations offer these levels.

The RAID Advisory Board

The RAID advisory Board is a consortium of manufacturers of RAID equipment and other interested parties. RAB is responsible for developing and maintaining RAID standards and has formal programs covering education, standardization, and certification. Supporting these three programs are six committees; Functional Test, performance Test RAID Ready Drive, Host Interface, RAID enclosure, and Education.

RAB sells several document, the most popular of which is The RAID book, first published in 1993. The RAID book covers the fundamentals of RAID and defines each RAID level. It's a worthwhile acquisition if you want to learn more about raid. Computer technology Review magazine hosts a RAB section in each monthly issue.

The RAB certification program awards logos to equipment that passes its compatibility and performance testing suites. The RAB conformance Logo certifies that a component so labeled complies with the named RAID level designation as published in The RAID book. The RAB Gold Certificate Logo certifies that a product meets that functional and performance specifications published by RAB. In mid 1996, 20 firms were authorized to apply the RAB Gold Certificate Logo to their products.

RAID 0

RAID 0 is high-performance zero-redundancy array option. RAID 0 isn't properly RAID at all. It stripes blocks of data across multiple disk drives to increase the throughput of the disk sub-system but offers no redundancy. If one drive fails in a RAID 0 array the data on all drives on the array is inaccessible. RAID 0 is used primarily for application needing the highest possible reading and writing data rate.

Nevertheless, there's place for RAID 0. Understanding RAID 0 is important because the same striping mechanism used in RAID 0 is used to increase performance in other RAID levels. RAID 1 is inexpensive to implement for two reasons:

- No disk space is used to store parity information eliminating the need to buy either larger disk drives or more of them for a given amount of storage.
- The algorithms used by RAID 0 are simple ones that don't add many overheads or require a dedicated processor.

RAID 0 uses striping to store data. Striping means that data blocks are alternately written to the different physical disk drives that make up the logical volume represented by the array. For instance, you RAID 0 array might comprise three physical disk drives that are visible to the operating system as one logical volume. Suppose that your block size is 8k and that a 32k file is to be written to disk. With RAID 0, the first 8k block again to drive 1. your single 32k file is thus sorted as four separate blocks residing on three separate physical hard disk drives.

This block-wise distribution of data across multiple physical hard disks introduces two parameters used to quantify a RAID 0 array. The size of the block used in this case, 8k us referred to as the chunk size, which determines how much data is written to a disk drive in each operation. The number of physical hard disk drive comprising the array determines the stripe width. Both chunk size and stripe width affect the performance of a RAID 0 array.

When a logical read request is made to the RAID 0 array (fulfillment of which requires that an amount of data larger than the chunk size be retrieved), this request is broken into multiple smaller physical read requests. Each read request is directed to and service by the individual physical drives on which the multiple blocks are stored. Although these multiple read requests are generated serially, doing so takes little time. The bulk of the time needed to fulfill the read requests is used to transfer the data itself. With sequential reads, which involve little drive head seeking, the bottleneck becomes the internal transfer rate of the drives themselves. Striping lets this transfer activity occur in parallel on the individual disk drives that make up the array, so the elapsed time until the read request is completely fulfilled is greatly reduced.

Striping doesn't come without cost in processing overhead, and this is where chunk size affects performance. Against the benefit of having multiple spindles at work to service

a single logical read request, you must weigh the overhead processing cost required to write and then read this data from many disks rather than just one. Each SCSI disk access requires numerous SCSI commands to be generated and then executed, and striping the data across several physical drives multiplies the effort required accordingly.

Reducing the block size too far can cause the performance benefits of using multiple spindles to be swamped by the increased time needed to generate and execute additional SCSI commands. You can actually decrease performance by using too small a block size. The break-even point is determined by your SCSI host adapter and by the characteristics of the SCSI hard disk drives themselves, but, generally speaking, a block size smaller than 8k risks performance degradation. Using block sizes of 16k, 32k, or larger offers correspondingly greater performance benefits.

Sequential reads and writes make up a small of total disk activity on a typical server disk subsystem most disk accesses are random; by definition, this means that you'll probably need to move the heads to retrieve a particular block of data. Because head positioning is a physical process, relatively speaking it's very slow. The benefit of striping in allowing parallel data transfer from multiple spindles is much less significant in random access because all the system components are awaiting relatively slow head positioning to occur. Therefore, striping does little to benefit any particular random-access disk transaction, but it does benefit random access disk throughput as a whole.

RAID 1

What do you do to make sure that you don't suffer by losing something? The obvious answer is to make a copy of it. RAID 1 works this way, making two complete copies of everything to mirrored or duplexed pairs of disk drives. This 100 percent redundancy means that if you lose a drive in a RAID 1 array, you have another drive with an exact duplicate of the failed drive's contents. RAID 1 offers the greatest level of redundancy, but at the highest cost for disk drives.

Mirroring means that each disk drive has a twin. Anything written to one drive is also written to the second drive simultaneously. Mirroring is 100 percent duplication of your drives. If one drive fails, its twin can replace it without loss of data.

Mirroring has two disadvantages:

- The most obvious is that you must buy twice as many disk drives to yield a given amount of storage.

- The process of writing to both drives and maintaining coherency of their contents introduces overhead, which slows writes.

Mirroring has two advantages:

- Your data is safely duplicated on two physical devices, making catastrophic data loss much less likely.
- Read performance is greatly increased because reads can be made by the drive whose heads happen to be closest to the requested data.

Duplexing is similar to mirroring, but it adds a second host adapter to control the second drive or set of drives. The only disadvantage of Duplexing, relative to mirroring, is the cost of the second host adapter although duplex host adapters, such as the Adaptec AHA-3940UW, are however, eliminates the host adapter as a single host adapters. Duplexing with two cards, however, eliminates the host adapter as a single point of failure.

In a large server environment the cost of duplicating every disk drive quickly adds up, making RAID 1 very expensive. With smaller servers, however, the economics can be very different. If your server has only one SCSI hard disk drive installed, you might find that you can implement RAID 1 for only the relatively small cost of buying another similar disk drive. At this writing, the cost of 1G of high performance SCSI –2 storage was about \$250, based on the \$1,150 street price of a 4.3 G Seagate ST3437W drive.

RAID 2

RAID 2 distributes the data across multiple drives at the bit level. RAID 2 uses multiple dedicated disks to store parity information and, thus requires that an array contain a relatively large number of individual disk drives. For example, a RAID 2 array with four data drives requires three dedicated parity drives. RAID 2 has the highest redundancy of any of the parity oriented RAID schemes.

The bit-wise orientation of RAID 2 means that every disk access occurs in parallel. RAID2 is optimized for applications such as imaging, which requires the transfer of large amounts of contiguous data.

RAID 2 isn't a good choice for random access applications, which require frequent small reads and writes. The amount of processing overhead needed to

fragment and reassemble data makes RAID 2 slow, relative to other RAID levels. The large number of dedicated parity drives required makes RAID 2 expensive. Because nearly all PC LAN environments have heavy random-disk access, RAID 2 has no place in a PC LAN. However, RAID 2 does have some specific advantages for special-purpose digital video servers.

RAID 3

RAID 3 stripes data across drives, usually at the byte level, although bit-level implementations are possible. RAID 3 dedicates one drive in the array to storing parity information.

Like RAID 2, RAID 3 is optimized for long sequential disk accesses in applications such as imaging and digital video storage, and is inappropriate for random access environment such as PC LANs. Any single drive in a RAID 3 array can fail without causing data loss, because the data can be reconstructed from the remaining drives. RAID 3 is sometimes offered as an option on PC-based RAID controllers but is seldom used. One primary application for RAID 3 is digital video and audio data storage for automated commercial insertion, direct-to-home satellite, and other TV broadcast uses.

RAID 3 can be considered an extension of RAID 0 in that it stripes small chunks of data across multiple physical drives. In a RAID 3 array that comprises four physical drives, for example, the first block is written to the first physical drive, the second block to the second drive, and the third block to the third drive. The fourth block isn't written to the fourth drive, however, it's written to the first drive to begin the round robin again.

The fourth drive isn't used directly to store user data. Instead, the fourth drive stores the results of parity calculations performed on the data written to the first three drives. This small chunk striping provides good performance on large amounts of data, because all three data drives operate in parallel. The fourth, or parity, drive provides the redundancy to ensure that the loss of any one drive doesn't cause the array to lose data.

For sequential data transfer, RAID 3 offers high performance due to striping, and low cost due to its reliance on a single parity drive. Its this single parity drive, however, that is the downfall of RAID 3 for most PC LAN applications. By definition no read to a RAID 3 array requires that the parity drive be accessed unless data corruption has occurred on one or more of the data drives. Reads, therefore, proceed quickly. However, every write to a RAID 3 array requires that the single parity drive be accessed and written to in order to store the parity information for the data write that just occurred. The random access typical of a PC LAN environment means that the parity drive in a RAID 3 array is overused, with long queues for pending writes, whereas the

data drives are underused because they can't proceed until parity information is written to the dedicated parity drive.

RAID 4

RAID 4 is similar to RAID 3, except RAID 4 strips data at the block or sector level rather than at the byte level, thereby providing better read performance than RAID 3 for small random reads. The small chunk size of RAID 3 means that every read requires participation from every disk in the array. The disks in a RAID 3 array are, therefore, referred to as being synchronized or coupled. The larger chunk size used in RAID 4 means that small random reads can be completed by accessing only a single disk drive instead of all data drives. RAID 4 drives are, therefore, referred to as being unsynchronized or decoupled.

Like RAID 3, RAID 4 suffers from having a single, dedicated parity disk that must be accessed for every write. RAID 4 has all the drawbacks of RAID 3 and doesn't have the performance advantage of RAID 3 on large read transactions. About the only environment for which RAID 4 would make any sense at all is one in which nearly 100 percent of disk activity is small random reads. Because this situation isn't seen in real world server environment, don't consider using RAID 4 for you PC LAN.

RAID 5

RAID 5 is the most common RAID level used in PC LAN environments. RAID 5 stripes both user and parity data across all the drives in the array, consuming the equivalent of one drive for parity information.

With RAID 5 all drives are the same size, and one drive is unavailable to the operating system. For example, in a RAID 5 array with three 1 G drives, the equivalent of one of those drives is used for parity leaving 2G visible to the operating system. If you add a fourth 1G drives to the array, the equivalent of one drive is still used for parity, leaving 3G visible to the operating system.

RAID 5 is optimized for transaction-processing activity, in which users frequently read and write relatively small amounts of data. It's the best RAID level for nearly any PC LAN environment and is particularly well suited for database servers.

The single most important weakness of RAID levels 2 through 4 is that they dedicate a single physical disk drive to parity information. Reads don't require accessing the parity drives, so they aren't degraded. This parity drive must be accessed for each write to the array, however, so RAID levels 2 through 4 don't allow parallel writes. RAID 5

eliminates this bottleneck by striping the parity data onto all physical drives in the array, thereby allowing both parallel reads and writes.

Windows NT Server RAID Options

On the reasonable assumption that software RAID is better than no RAID, using Windows NT server 4.0 to provide RAID functionality makes sense, particularly for small servers that other wise wouldn't be equipped with RAID functionality. The following RAID options available with Windows NT server.

RAID 0 is referred to by Microsoft as Disk Striping, or the use of stripe sets. Like any RAID 0 arrangement stripe sets increase disk subsystem performance but do nothing to provide data redundancy. The failure of any disk drive in a stripe set renders the data on the remaining drives in the stripe set inaccessible. Windows NT server allows a stripe set to comprise from 2 to 32 individual disks. Increasing the number of disks in a stripe set increases the probability of data loss, because failure of a single drive results in a failure of the entire stripe set. If the disks assigned to a stripe set vary in size, the smallest determines the common partition size of the stripe set. The remaining space on other drives in the stripe set might be used individually or assigned to a volume set.

RAID 1, referred to by Microsoft as Disk Mirroring, or the use of mirror sets, is supported directly by windows NT server for any hardware configuration with at least two disk drives of similar size. Windows NT server doesn't require that the mirrored drive be identical to the original drive, but only that the mirrored drive be at least as large. This considerably simplifies replacing failed drives if the original model is no longer available. RAID 1 duplexing is supported directly for any hardware configuration with at least two disk drives of similar size controller controllers. As with any duplex arrangement, this removes the disk controller as single point of failure. As with mirrored drives, windows NT server doesn't require duplexed drives to be identical.

RAID 5 referred to by Microsoft as Disk Striping with parity is also supported natively by windows NT server for any hardware configuration with at least three disk drives and one or more disk controllers. Windows NT server allows as many as drives in a striping set, although to maintain

acceptable performance when a single drive fails, it is a better idea to limit the RAID 5 array to five or six drives.

13.3 Creating Windows NT Server Stripe and Mirror Sets

After you install a sufficient of drives and ensure that your SCSI host adapter recognizes each drive, you can create one or more of the three levels of RAID supported by windows NT server 4.0. The following sections provide the instructions for implementing windows NT 4.0 software RAID 0,1, and 5, RAID 0 and 1 require two physical drives; RAID 5 requires three physical drives. Each drive must have unused space in which to create the stripe set or mirror volume.

Creating RAID 0 Stripe Sets

Creating a RAID 0-stripe set is the simplest of the three processes offered by Windows NT 4.0's Disk Administrator. To create a RAID 0 stripe set, proceed as follows.

1. Log on as administrator and run Disk Administrator by choosing programs, administrative Tools, and Disk Administrator from the Start menu.
2. Click to select an unused area on the disk drive (See the figure)
3. Ctrl+click to select additional unused areas on other disk drives, up to a total of as many as 32 disk drives.
4. From the partition menu, choose create strips set to open the create stripe set dialog, which displays in the create stripe set of total size box the default total size of the stripe set spanning all selected drives. This total size takes into account the smallest area selected on any disk drive, adjusting the sizes of the areas on the other selected drives to set all to identical size. The total size value is approximately the size of the smallest disk area multiplies by the number of drives in the stripe set.
5. Click OK to accept the default size. Windows NT server prepares to create the stripe set and assigns it a single default drive letter. At this point, no changes have been made to the drives.

6. From the partition menu, choose commit changes now. A message box informs you that changes have been made to your disk configuration. Click yes to accept and save the changes.
7. A second message box notifies you that the update was successful, and that you should save the disk configuration information and create a new emergency repair disk. Click OK.
8. Click to select the newly created but unformatted stripe set. From the Tools menu, choose format to open to format drive H:
9. Type a name for the volume in the label text box. In the format type drop-down list, you can choose and NTFS or FAT file system. Select NTFS (the default). Marketing the Quick Format check box bypasses drive sector checking during the formatting process. Click OK to continue.
10. A confirmation dialog warns you that continuing overwrites the contents of the volume. Click yes to continue formatting.
11. If your marked the Quick Format check box in the format drive H: dialog, the quick format progress indicator appears only briefly. For a conventional formatting operation, the progress indicator appears for a minute or more, depending on the volume.
12. When formatting is complete, a dialog informs you of the total available disk space on the new volume. Click OK. The new volume is ready for use.

Creating drive configuration and Emergency Repair Disks

After making permanent changes to your drive configuration, always save the configuration changes, replace there pair information on the fixed disk, and create a new emergency repair disk. Follow these steps to ensure that you can restore your existing configuration in case of a system failure.

1. To save the current drive configuration, choose configuration and then save from the partition menu to display the insert disk dialog. Insert a formatted disk and click OK to save the configuration. In case of major castastrophe, you can use the disk to restore the current drive configuration by choosing configuration and then restore from the partition menu.

2. Choose Run from the Start menu, type rdisk in the open text box, and click OK to open the repair disk utility window
3. Click the Update Repair Info button. The confirmation dialog appears. Click OK to continue the update, which takes a minute or two.
4. After the local repair information is updated, a message box appears, asking whether you want to create an emergency repair disk. After updating local repair information it's imperative that you create a new emergency repair disk. Insert a disk in the disk drive and click yes.
5. Replace your existing emergency repair disk with the new emergency repair disk, which should be stored in a safe location. Consider making two emergency repair disks, storing one disk off site.) The old disk is unusable with the new repair information stored on the local fixed disk.

Creating RAID 1 Mirror Sets

Mirror sets vary from strips sets; whereas strips sets may span as many as 32 drives, mirror sets are created on a paired drive basis. You must first create a standard formatted volume, and then creates the mirror drive.

Creating and formatting a new standard volume To create and format new separate volume from the free space available on a single drive ,follow these steps.

1. Log on as administrator and run Disk Administrator by choosing programs, administrative, Tools, and Disk administrator from the Start menu.
2. Click to select an unused area on the fixed disk drive.
3. From the partition menu, choose create to open the create logical drive dialog, which displays in the create logical drive of size box the default total size of the free space of the selected drive. Accept the default size unless you want to creates a volume of a smaller size.

4. Click OK. Windows NT server prepares to create the stripe set, assigning the volume a single default drive letter. At this point, no changes have been made to the drive and the drive is disabled.
5. From the partition menu, choose commit changes now. A message box informs you that changes have been made to your disk configuration. Click yes to accept and save the changes. A second message box notifies you that the update was successful, and that you should save the disk configuration information and create a new emergency repair disk. These two steps are performed after you complete the drive reconfiguration process. Click OK.
6. Click to select the newly created but unformatted volume. From the Tools menu, choose Format to open the Format Drive J: dialog:
7. Type a name for the volume in the Label text box. From the Format Type drop down list you can choose an NTFS or FAT file system; select NTFS. Marking the Quick Format check box bypasses drive sector checking during the formatting process. Click OK to continue.
8. A confirmation dialog warns you that continuing overwrites the contents of the volume. Click yes to continue formatting. If you marked the Quick Format check box, the Quick Format progress indicator appears only briefly.
9. When formatting is complete, a dialog informs you of the total available disk space on the new volume. Click OK. The new volume is ready for use as an independent volume or as a member of a mirror set.

Creating the Mirror of the Standard Volume

Mirroring creates a formatted volume of the same size as the new standard volume, but on another physical drive. To create the mirror partition, follow these steps;

1. From Disk Administrator, select the newly formatted standard volume.
2. Press Ctrl click an unused area on another disk drive that's at least as large as the newly created volume.
3. From the Fault Tolerance menu, choose Establish Mirror.

4. From the Partition menu, choose now. Windows NT server creates the mirror set and assigns the drive letter of the first drive of the set.
5. Update the configuration and repair information as described earlier in the section “Creating Drive Configuration and Emergency Repair Disks”.

Creating RAID 5 Stripe Sets with Parity

The process of creating a stripe set with parity is very similar to that used to create a RAID 0 stripe set. Whereas a RAID 0 stripe set can be created on only two physical drives, a RAID 5 stripe set with parity requires a minimum of three drives one for parity information and at least two for data.

To create a stripe set with parity proceed as follows;

1. Log on as administrator and run Disk Administrator by choosing Programs, Administrative Tools, and then Disk administrator from the start menu.
2. Click to select an unused area on the first disk drive.
3. Press the Ctrl button and click at least two additional unused areas on other disk drives, up to total of as many as 32 disk drives. You can choose only one area on each disk drive.
4. From the Fault Tolerance menu, choose Create Stripe Set with Parity. The Create Stripe Set with Parity dialog appears, displaying the total size of the stripe set with parity spanning all selected drives. This total size takes into account the smallest area selected on any disk drive, adjusting the sizes of the areas on the other selected drives to set all to identical size.
5. Click Ok to accept the default. Windows NT server prepares to create the stripe set with parity and assigns a drive letter. The stripe set with parity must now be prepared for use.
6. From the partition menu, choose commit changes Now. A dialog tells you that changes have been made to your disk configuration. Click Yes to accept and save the changes. A message box notifies you that the update was successful. Click OK.

7. Select the newly created but unformatted stripe set with parity. From the Tools menu, choose Format to open the Format Drive H:(may be any alphabet) dialog;
8. Type a name for the volume in the Label text box. From the Format Type drop-down list you can choose an NTFS or FAT file system, select NTFS. The Quick format check box is disabled when you create a RAID 5 stripe set. Click OK to continue.
9. When formatting is complete, a dialog informs you of the total available disk space on the new volume. Click OK and the new RAID 5 volume is ready for use.

Saving the Master Boot Record and Partition Table

After you finish setting up your software RAID subsystem, use the DiskSave or DiskProbe utilities of the Windows NT server Resource Kit to create a copy of the master boot record and partition table on a disk. You might need to use the saved file in the event of your inability to recover from a drive failure, especially with a RAID 1 mirror set.

13.4 Recovering a Software RAID 1 or RAID 5 Set

Failure of a member of a mirror or stripe set with parity results in the creation of an orphan volume. Windows NT 4.0 detects an orphan volume under either of these two conditions:

- During a failed write operation, in which case you receive an error message.
- During Windows NT startup, in which case a server error is logged in the event log. In either case, operations continue, taking advantage of Windows NTs RAID 1 or RAID 5 fault tolerance. Disk administrator's status line displays volume #N (recoverable) for the selected volume. When a drive fails, create a backup with a fresh tape immediately your RAID set is no longer fault tolerant, and failure of another drive will result in loss of all data since the last backup.

Restoring a Mirror Set

To recover a failed RAID 1 mirror set, you must first break the mirror set, and then reconstruct it after installing a new drive, preferably of the same make and model as the failed drive. After making a complete backup of the remaining good drive(s) of the mirror set, follow these steps to recover a mirror set.

1. Launch Disk Administrator and select the mirror set volume.
2. Choose break mirror from the fault tolerance menu. You must break the mirror set to recover.
3. Click Yes to confirm that you want to break the mirror set. The working partition receives that logical drive letter assigned to the set, and the orphan partition receives the next available drive letter.
4. Choose Commit changes now from the Partition menu.
5. Make an independent backup of the working partition.
6. Shut down the system and remove the failed drive. If the failed drive is the shadow drive, set the device ID of the shadow drive to that of the primary drive. Replace the failed drive with the same make and model set to the appropriate device ID.
7. Restart the system.
8. Re-create the mirror set by using the procedure described earlier in the section "Creating the Mirror of the standard volume.
9. Re-create your emergency repair disk and use the Resource Kit's DiskSave or DiskProbe utility to save the master boot record and partition table to disk.

Restoring a Stripe Set with Parity

Restoring a RAID 5 subsystem is much easier than recovering a RAID 1 mirror set. To recover from failure of a RAID 5 drive, follow these steps:

- 1 After making a full backup shutdown the system and replace the failed drive.
- 2 Restart the system, launch Disk Administrator and select the RAID 5 volume
- 3 Select the unpartitioned space on the replacement drive.

- 4 Choose Regenerate from the Fault Tolerance menu
- 5 Choose Disk Administrator and restart the system . Regeneration of the failed drive's contents occurs automatically in the background after Windows NT restarts.
- 6 Re-create your emergency repair disk and use the Resource Kit's DiskSave or DiskProbe utility to save the master boot record and partition table to disk.

13.5 Short Summary

- A RAID system combines two or more disks to create a large virtual disk structure that enables you to store redundant copies of the data.
- Disk Duplexing uses two controller cards, one for each disk. Disk Duplexing is the same as disk mirroring but disk mirroring uses only one controller card
- Duplexing protects information at the hardware level with duplicate channels and duplicate hard drives
- Striping without parity should be reserved for workstations and servers on which speed considerations are paramount and possible downtime is an acceptable risk.
- RAID 5 requires at least three drives, because this version writes data across two of them and then creates the parity block on the third.
- Striping with parity is based on the principle that all data is written to the hard drive in binary code (zeros and ones).
- RAID 0 uses data striping and block interleaving, a process that involves distributing the data block by block across the disk array in the same location across each disk.

13.6 Brain Storm

1. What are the RAID levels, which is available in the market?
2. Describe the difference between Disk Striping and Mirroring?
3. Give a notes on Disk Duplexing?
4. How does RAID Level 1 operate?
5. What is the minimum number of disks needed to configure a stripe set with parity on Windows NT Server?
6. What is the maximum number of drives supported in a mirror set?
7. What method of fault tolerance uses two controller cards on two separate hard drives?



Lecture 14

Installing Backup Filling Systems

Objectives

In this Lecture you will learn the following:

- Understanding the Backup Types
- Able to choose the Backup Hardware
- Knowing the NT Server 4.0 Backup Application

Coverage Plan

Lecture 14

14.1 Snap Shot

14.2 Backup Types

14.3 Developing a Backup Strategy

14.4 Choosing Backup Hardware

14.5 Windows NT Server 4.0 Backup Application

14.6 Short Summary

14.7 Brain Storm

14.1 Snap Shot

In this lecture you will understand the types of backup available and how each type of backup can best be used to ensure the safety of your data, Backup application and how to backing up open files.

Backing up your data to tape or other removable storage media serves these purposes:

- You have a copy of your data to protect against the catastrophic data loss that occurs when a disk drive fails and your RAID subsystem can't regenerate the data
- You get an offline copy of data that you can recover if the working copy of a file is deleted or improperly modified
- You get a data archive that can be preserved for historical or legal purposes
- You can maintain a copy of your data, off premises to protect against fire or other natural disasters.

14.2 Backup Types

Windows NT server 4.0 provides built in support for various backup devices, primarily Tape backup drives. The Microsoft hardware compatibility list for Windows NT, current when this book was written, identifies 145 individual makes and models of backup devices compatible with Windows NT server 4.0. Before you choose backup hardware and software for your windows NT server 4.0 installation, however, it is important to understand how windows NT server 4.0 handles the files backup process. The following sections explains the purpose of the file archive bit and the types of backup operations supported by windows NT server 4.0 and third party device drivers for backup devices.

The Archive Bit

For you to manage a backup strategy, it is essential that your backup software has a way of knowing when a file has been created or modified since the last normal backup. One way to do this is to examine the date/ time stamp on each file, and compare it with the time that the last backup was done to determine whether the file has changed. This method, used by the backup applet bundled with windows 95 is simple in concept but unreliable in use. Using data/time stamps to determine the files to back up is unreliable, because many programs alter the contents of a file but don't change the date/time when doing so. As a result, a better method of determine backup currency is need. Fortunately, such a method has been available since the early days of MS-DOS.

Normal Backups

A normal backup copies all selected files, regardless of the state of their archive bits, to the tape drive or other backup media, and then runs off the archive bit on all files that have been copied. Most third party backup software refers to this process as a full backup.

Recreating a failed hard drive from a normal backup set of the entire drive is straightforward. If the system drive has failed, you must replace the drive and reinstall windows NT server before proceeding. If the system drive is operable, after replacing the data drive, use the Windows NT server backup application. NTBACKUP, to do a full restore of the tape to the new drive. Partial restores, such as those of accidentally deleted files, are equally straightforward. NTBACKUP is described later in the section “using the windows NT server 4.0 Backup Application.

Copy Backups

A copy backup is identical to a normal backup, except that copy backups skip the final step of resetting to off all the archive bits on backed up files. Most third party backup software refers to this process as a full copy backup. The resulting backup tape is identical to what would have been created by a normal backup, but the archive bit status of the files on the disk remains unchanged. The main purpose of a copy backup is to allow you to create an archive or off site backup set without affecting your main backup set’s rotation process.

Because the contents of a copy backup set are indistinguishable from those of a normal backup set, restore procedure are identical for these two types of backups.

Incremental Backups

An incremental backup copies to the backup media all selected files that have their archive bit turned on, and then turns the archive bit off for the files that have been copied. The tape from the first incremental backup done after a normal backup contains only those files altered since the last normal backup. Subsequent incremental backup tapes contains only those files that changed since the last incremental backup. After each incremental backup is completed, all files have their archive bits turned off as though a normal backup were done.

Re-creating a failed hard drive from incremental backup sets is a bit more involved than using a normal backup set, because each incremental backup tape contains only some of the changed files, and different incremental backup tapes might contain different versions of the same file.

To re-create a failed drive, you first restore the most recent normal backup set to the replacement disk. Then you restore all incremental backup sets created after the normal backup set, beginning with the earliest and proceeding sequentially to the latest incremental backup set.

Restoring an accidentally deleted file is a more complex process. To ensure that you get the latest version of the file, you must start by examining the most recent incremental backup set and work backward until you locate the most recent occurrence of file on an incremental backup set. If the file in question hasn't changed since the last normal backup, you might have to work all the way back to the last normal backup set before you locate the file. Fortunately, most backup software makes this process somewhat easier by allowing you to search backup logs to locate the file so that you can load the proper tape directly.

Differential backups

A differential backup copies to the backup media all selected files that have their archive bits turned on, but then leaves the archive bits unchanged on the files that have been copied. Each differential backup set contains all files changed since the last normal backup. Each differential backup set also is larger than the preceding set, because later sets contain all the files previously backed up, plus all files changed since that last backup.

Re-creating a failed hard drive with a differential backup set is relatively straightforward. As with an incremental backup set, you begin by restoring the last normal backup set. Because each differential backup set contains all files changed since the last normal backup, however, you need to restore only the most recent differential backup set. Restoring an accidentally deleted file is similarly straightforward. If the file is listed on your most recent differential backup log, restore from the latest differential backup set. Otherwise, restore from the last normal backup.

Daily copy backups

A daily copy backup copies all selected files that have been modified that day, but leaves the archive bits unchanged on the copied files. Like the windows 95 backup mentioned earlier in the section "The Archive Bit," the daily copy backup uses file date stamps to determine their eligibility for backup, rather than examine the status of the archive bit.

Unlike the incremental backup, which copies all files changed since the last normal backup or incremental backup was done, the daily copy backup must be run at least

once each day if it's to be successfully used to archive files changed since the last normal backup. If you fail to run the daily copy backup on one particular day, none of the files changed on that day are written to tape until the normal, copy, incremental, or different backup is done. Because it ignores the state of the archive, but the daily copy backup also fails to back up changed files if the file date stamp wasn't altered at the time the file was changed.

14.3 Choosing Backup Hardware

Backup and tape drives traditionally are considered synonymous. Although tape is the overwhelming choice of backup media for most network administrators, various optical storage technologies are beginning to nip at its heels. These technologies are still niche products, in so far as the backup market is concerned, because they are largely proprietary in nature and because their cost per byte stored is still relatively high in most cases. Still it is worth examining some of these alternatives briefly for their current value in fitting specific needs and their possible future values as an alternative to tape. The following sections describe the relative merits of common tape backup formats and alternative optical storage systems for backup and archiving data.

Tape Drives and Formats

Tape drives are the traditional method for backing up data. In terms of reusability and cost per byte stored, the tape drive is now and likely will remain the best choice for backing up windows NT 4.0 servers. The following sections describe the most common types of tape drives in use today.

Quarter-inch cartridge (QIC) drives

Most tape drives sold today are quarter inch cartridge or QIC compatible. These factors are used for QIC cartridges:

- Full – size cartridges are about the size of a small paperback book and are commonly referred to as a DC-600, although various versions of this cartridge exist.
- Smaller cartridges are the familiar 3 ½ inch mini cartridge commonly called a DC-2000, but many versions of this cartridge exist.

Digital Linear Tape

The high end tape backup hardware and media is called Digital Linear Tape(DLT). Because it's very fast, offers large capacities, and is extremely reliable, DLT is beginning to replace DAT and 8mm tape drives in large server environments. Current

DLT drives have capacities of 40 compressed, or 20G uncompressed so even a very large disk subsystem can be backed up to a single tape.

Helical Scan Tapes

Originally developed for use in video recording, helical scanning works by running a tape past a head that rotates at an angle relative to the motion of the tape. The resulting tracks resemble diagonal lines running from one edge of the tape to the other, repeating this pattern from end to end on the tape.

Writable optical drives

Writable optical drives today are used primarily for archiving data, rather than backup. Various new technologies include erasable CD-ROMs (CD-E) and writable digital video discs (DVDs, also called digital versatile discs) which store up to 4.7G. Even farther in the future are exotic writable disc technologies, based on cholesteric liquid crystals which promise up to 280G of storage per side.

The following sections describe currently available drives that use lasers to write and read from discs that, for the most part, have the same dimensions as conventional audio CDs and CD-ROMs.

CD-Recordable Drives

Recordable CD(CD-R) technology has been around for a few years, but it's just now joining the mainstream. With prices on drives dropping below the magic \$ 500 point, sales of CD-R drives are climbing rapidly. CD-R drives are similar to standard CD-ROM drives, but CD-R drives use a higher powered laser that can write to specially designed CDs. These CDs can then be read in any standard CD-ROM drive.

CD-R's relatively low capacity of 680M, its use of relatively expensive media (about \$5 per disc), and its lack of rewritability make CD-R a poor choice for routine backup. The first two issues are likely to be addressed at least incrementally as the technology improves, although revolutionary improvements are unlikely. The read only nature of a CD-R disc can be an advantage for applications such as data archiving.

Magneto-Optical Drives

Magneto-optical(MO) disks are another technology sometimes considered for use as a backup media. Magneto-optical disks use a combination of a high power laser and a magnetic head to write to their media. The laser heats the media, allowing the magnetic

head to realign the magnetic particles. Because this action is repeatable, MO disks are read write like a traditional disk drive, rather than write once like CD-R and WORM drives.

MO drives now have performance more similar to that of hard disk drives than the performance usually associated with optical drives. However, Mo drives have relatively low capacity and high media costs, making them inappropriate as backup devices for most situations.

Worm Drives

Write once, read many technology has been available longer than either CD-R or MO devices. WORM drives are available in various platter sizes up to 12 inches and in capacities of up to 6G per disc. WORM jukeboxes can provide near online storage capacity in the terabyte range. WORM drives are incrementally rewritable (data can be added incrementally) allowing backup of multiple versions of the same file or folder to a WORM disk.

WORM is an excellent if expensive archiving medium. For applications that require storing huge amounts of data, such as document imaging, the nearly online performance of WORM can be considered adequate for online use.

14.4 Windows NT Server Backup Application

The windows NT server 4.0 backup applications, NTBACKUP, has two obvious advantages it's included with windows NT 4.0 and, as a bundled application compatibility and reliability problems are less likely to occur. Balanced against these advantages are NTBACKUP's paucity of high end features and limited options. If your LAN relatively small, the LAN's architecture is simple, and your backup requirement are modest, NTBACKUP suffices. For single server environments that use a simple tape rotation method and have no need to backup workstations from the server, NTBACKUP is more than adequate.

NTBACKUP can back up files stored on a drive that uses either the NTFS or FAT file systems, and it can restore the files backed up from a drive that uses one file system to a drive that uses the other file system. NTBACKUP does only file by file backups and makes no provision for doing a disk image backup. It supports only tape drives as destination devices. You can't back up from one hard drive to another by using NTBACKUP.

Setting up NTBACKUP for use with your tape drive

The NTBACKUP program files are installed when you install windows NT server. Before using NTBACKUP however, you must first install support for your tape drive by following these steps.

1. Form control panel, double click the Tape Devices tool to display the Tape Devices property sheet. The Devices page shows installed tape devices. You highlight a displayed tape device and click the properties button to display the properties for that device. You also can click the Detect button to attempt to automatically detect and install a driver for a physically installed tape device that's not shown in the list.
2. To install a new tape device manually, click the Driver tab. Windows NT builds a driver list for several seconds, and then displays the drivers page. The drivers page shows a list of installed tape device drivers.
3. You can click the Remove button to remove a currently installed driver. If no driver is shown for the tape drive you want to install, click the Add button. Windows NT 4.0 again creates a driver list. After a few seconds, the Install Driver dialog appears
4. Select the manufacturer of your tape drive in the manufacturers list and then select one of the supported tape devices made by that manufacturer from the Tape Devices list. Click OK to install the driver for that tape drive, or cancel to abort the process.
5. You're prompted to insert the windows NT 4.0 distribution CD-ROM disk into your CD-ROM drive. You can type the path where the files are located, select the location from the drop down list, or browse for the proper location. The distribution files for a server running an Intel processor are located in the \1386 folder of your CD-Rom drive.
6. After the Copy Files From text Box is completed properly, click OK to begin installing the new tape device driver.
7. When the files are copied, you're prompted to restart windows NT for the changes to take effect.

Backing up with NTBACKUP

Using NTBACKUP to back up your files requires several steps. You must prepare and label the select the volumes, folders and files to be included in the backup set; choose the appropriate backup options to use and then run the backup itself. Optionally, you may choose to run a comparison pass after completing the backup to verify the integrity of your backup set.

Preparing and Labeling Media

Depending on the type of tape drive you use, there might be little or quite a lot of preparation needed. QIC tape drives require that tapes be formatted before use and that they be periodically retensioned to avoid breakage. The QIC formatting process takes considerable time, but you can avoid such time consumption by purchasing preformatted tapes. Like VCR tapes (the other technology that uses helical scanning) DAT and 8mm tapes require neither formatting before use nor periodic retensioning.

NTBACKUP gives you several tape tools. NTBACKUP allows you to erase tapes by using either a standard erase (which simply deletes the header information) or a secure erase (which actually overwrites the data on the tape). NTACKUP allows you to retension tapes that require such maintenance, and allows you to eject a tape if your tape drive supports software controlled ejection.(now there's useful feature)

Selecting Drives, Folders, and Files to Be Backed Up

After you prepare and label your tapes, the next step is to run NTBACKUP and select the drives, folders, and files to be backed up. Proceed as follows:

1. From the Start menu, choose Programs, Administrative Tools(Common), and Backup to open the Backup – [Tapes] window. If a blank formatted tape is in the tape drive, it is shown in the left pane as Blank Tape. If the tape contains data, the left pane displays the creation date and other particulars of the tape, and the right pane displays a brief summary of the tape's contents.
2. To begin selecting the drives, folders and files to be backed up, from the Window menu choose Drives. The Backup – [Drives] window appears(see Figure 14.1).

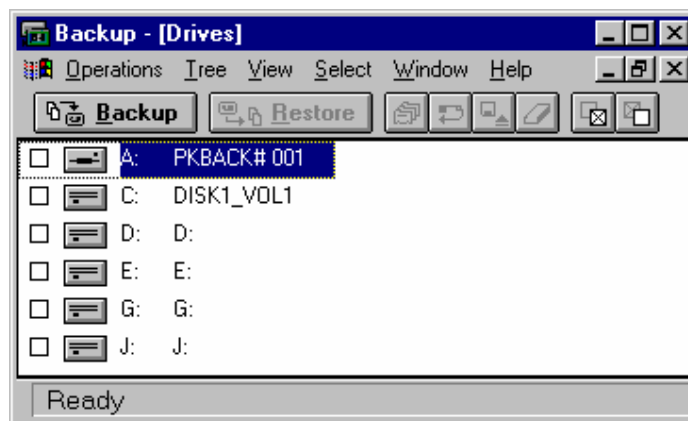


Figure 14.1 The Backup [Drives] window, displaying drives accessible to NTBACKUP

3. To select an entire drive to be backed up, mark the check box to the left of that drive's icon. To back up some of the folders and files on a particular drive, double-click the icon for that drive to view a folder tree for that drive. The folder tree view is similar to Windows NT Explorer, showing folders in the left pane and the files contained in a selected folder in the right pane. Single-clicking a folder name in the left pane displays the files contained in that folder in the right pane.
4. Select folders to be backed up by marking the check box to the left of the folder, which selects all files contained in that folder and its subfolders. Alternatively, you select individual files within a folder by marking the check box for that file or files in the right window pane and leaving the folder name in the left pane unmarked.
5. Repeat steps 3 and 4 for any other drives to be backed up.

Choosing Options and Running the Backup

After you select the drives, folders, and files to be backed up, the next step is to choose backup options. To do so, click the Backup button in the Backup window to display the Backup Information dialog. The Backup Information dialog is divided into three sections. The first section presents items about the currently mounted tape, including its name, owner, and creation date. The first section also allows you to specify various options that determine how the backup job is run.

The first set of backup options are as follows:

- The *Tape Name* text box allows you to enter a descriptive name for this tape.
- The *Verify After Backup* check box allows you to specify whether you want to do a second comparison pass after the backup is complete. Leaving this option selected ensures that the backup data is written to tape successfully, but nearly doubles the time needed to complete the backup.
- The *Backup Local Registry* check box lets you specify whether the Registry is backed up. You should back up the Registry as a matter of course. Don't clear this check box unless you have a very good reason for doing so.
- The *Operation* section lets you choose whether the current backup data is added to the end of the current tape or over writes the current contents of the tape. Append adds and Replace overwrites.
- The *Restrict Access to Owner or Administrator* check box lets you secure the contents of the tape against unauthorized access. When this check box is marked,

the contents of the tape can subsequently be accessed only by the tape's owner or by a member of the Administrators or Backup Operators group. This check box remains disabled (gray) unless you select the Replace option in the Operation section.

- The *Hardware Compression* check box determines whether hardware compression is used on drives that support it. If your drive doesn't provide hardware compression, this check box is disabled.

The second section of the Backup Information dialog, Backup Set Information, lets you specify information about particular backup sets. The scroll bar lets you view one such set at a time (you see the scroll bar only if you have selected more than one drive backed up). Each backup set includes the following information:

- *Drive Name* is the drive (for example, C) that you picked during the selection process described earlier.
- *Description* is a free text area in which you can enter a brief meaningful description of the contents of the backup set.
- *Backup Type* is a drop-down list from which you can select Normal, Copy, Incremental, Daily Copy, or Differential to do your backup.

The third section of the Backup Information dialog, Log Information, lets you specify logging information for the backup session, including the log file name and location, and the level of detail you want the log to contain.

- The *Log File* text box allows you to enter the drive, folder, and name of the file to which session information is to be logged as an ASCII text file.
- When selected, the *Full Detail* option button causes all transaction detail to be logged, including the names of all files and folders that are backed up. Although this level of detail can be useful at times, the resulting log files can be huge.
- When selected, the *Summary Only* option button logs only critical exception information, such as tape mounting information, start time of the backup, and files that were selected for backup but couldn't be opened.
- When selected, the *Don't Log* option button disables all logging functions.

After you select the options desired, click OK to begin the backup. The backup then begins, unless you've chosen the Replace option and the target tape already contains

data. If so, you'll be prompted to make sure that you really want to overwrite the tape. When the backup completes, a compare pass begins, if you chose that option.

As the backup begins, the Backup Status dialog appears, displaying the progress of your backup until it completes.

Restoring with NTBACKUP

Restoring with NTBACKUP is relatively straightforward. NTBACKUP allows a backup set to be restored to the same system from which it was made, or to a different system. NTBACKUP also allows backup sets made from either supported file system to be restored to a disk that uses NTFS or FAT. The following steps show you how to use NTBACKUP to restore files to your system:

1. Insert the tape that contains the data to do restored into your tape drive.
2. From the Start menu, choose Programs Administrative Tools(Common) and Backup to run NT Backup. The Backup [Tapes] window appears. The left pane of this windows displays the name and creation date of the tape. The right pane displays the folder backed up to the root level of the tape, the set number, and the number of the tape within that set.
3. Double-click the folder name in the right pane of the window. NTBACKUP loads the catalog from tape and displays the folder tree contained on the tape
4. Select the drives folders, and files to be restored in the same manner as you specified them earlier in the section "Selection Drives, Folders, and Files to Be Backed Up."
5. After you select all the drives, folders and files you want to restore, click the Restore button to begin restoring. The Restore Information dialog appears.

The Restore Information dialog is divided into two sections. The Backup Set Information section display properties of the mounted tape, including its name, the backup set of which it is a member, its creation date, and its owner. The Backup Set Information section also allows you to set the following options.

- ◆ The *Restore to Drive* drop-down list allows you to select the drive to which the backup set is to be restored. NTBACKUP can restore to any drive visible to Windows NT, including those on other servers.

- ◆ The *Alternate Path* text box allows you to restore to a path other than that from which the files were originally backed up.
- ◆ When selected, the *Restore Local Registry* check box causes NTBACKUP to restore Registry information if it's present on the backup tape.
- ◆ The *Restore File Permissions* check box determines whether ACL(Access Control List) information will be restored with the file. If this box is selected, the original ACL information is recreated as the file is restore. If the box is deselected the restored file instead inherits the ACL information of the directory to which it's resorted. If you're restoring to a FAT partition, this check box remains grayed out, because windows NT Server doesn't support permissions on file systems other than NTFS.
- ◆ When selected the *Verify After Restore* check boxes causes NTBACKUP to perform a comparison pass after completing the restore to verify that restored files correspond to those stored on tape. Exceptions are written to the log file.

The second section of the Restore Information dialog, Log information, allows you to specify logging information for the restore session, including the log file name and location, and the level of detail you want the log to contain;

- ❖ The *Log File* text box allows you to enter the drive, folder, and name of the file to which session information is to be logged as an ASCII text file. Illustrates a typical log file opened in Windows Notepad.
- ❖ The *Full Detail* option button, when selected causes all transaction detail to be logged.
- ❖ The *Summary Only* option button, when selected, logs only critical exception information for example, files selected for restore that couldn't be read from tape.
- ❖ The *Don't Log* option button, when selected disables all logging functions

After you select the options desired, click OK to begin the restore. The restore then begins and runs to completion, with a message that tells you the process is complete. When the restore completes, a compare pass begins(if you selected that option).

14.6 Short Summary

- A normal backup doesn't necessarily copy all files from a particular volumes or disk drive, but it might simply copy a file or set of files from a specified folder on the selected volume or disk drive.

- The incremental backup is best suited to environment where a relatively large number of different files change each day.
- Incremental backup sets the archive bit to off after each file is backed up, each file is backed up only on the day that it's changed. This might reduce the number of tapes needed for each daily tape set, and also cuts down on the time required for daily partial backups.

14.7 Brain Storm

1. What is the need to backup a file or folder?
2. What are backup types?
3. How you can restore the backup files?
4. Write the steps to install the tape device?
5. Define Tape drives and list the type of tape drives?



Lecture 15

Windows NT Register

Objectives

In this Lecture you will learn the following:

- Understanding the concept of Registry
- Able to configure the settings in Registry
- Able to maintain the registry security

Coverage Plan

Lecture 15

- 15.1 Snap Shot
- 15.2 Registry Basics
- 15.3 Configuration settings in Registry
- 15.4 Registry's Organization
- 15.5 Registry Editor
- 15.6 Important Hives & Keys
- 15.7 Inspecting Another Computer's Registry
- 15.8 Maintaining Registry Security
- 15.9 Short Summary
- 15.10 Brain Storm

15.1 Snap Shot

The Registry has been around since the beginning of Windows NT development. If you were a Windows 3.1 user or administrator, remember the chaos that occurred with INT files. These files were text files that held initialization information for Windows 3.1 and Windows application, unfortunately, the files were easy to modify, they easily corrupted, and all Windows application created their own INT files, filling the Windows directories with those files.

The Registry was the NT developer's answer to these and other problem. The information that was previously found in the INT files is now found in the Registry. In fact, it worked so well that the Windows 95 product group copied many of the concepts of the NT Registry.

15.2 Registry Basics

The registry is a binary database of the setting that Windows NT and its application need to start and operate. Because the Registry is a binary file, the only way to edit it is with the Registry Editor provided by Microsoft or with programming tools. As an administrator, focus on using the Registry Editor.

The Registry Editor isn't found on any of the default menus. Run the program to begin using the editor. In Windows 4.0, Microsoft released two versions of the Registry Editor. One version has the same functionality as the Windows 95 Registry Editor; the other has the same functionality as the Windows NT 3.51 Registry Editor. We'll explain on this topic later, but because of the limitations of the Windows 95 Registry Editors in a Windows NT environment, we'll focus on the Windows NT Registry Editor.

15.3 Configuration settings in Registry

The Registry keeps track of all configuration information for a computer, including applications, hardware, device drivers and network protocols. The Registry is a set of flat file databases each having a hierarchical organization. Windows NT uses the Registry during the boot process to determine which device drivers to load and in what sequence to load them. Registry entries store desktop settings for one or more users. Most of the administrative tools you use to configure Windows NT Server and keep it running smoothly alter. Registry entries. The following sections briefly describe how Windows NT uses the Registry.

Types of Registry information

The Registry is the one place that stores virtually everything that Windows NT needs to know about our hardware software and the users who may log on to the system. The Registry stores the following types of information:

- ❑ Information about your hardware and the device drivers required by the hardware
- ❑ A list of services such as SQL Server and Exchange Server, to start immediately after the boot process.
- ❑ Network information, including details about each network interface card and protocol in use.
- ❑ OLE and ActiveX information such as the file name and location of each OLE Server and ActiveX component
- ❑ File association information What application launches what kind of file and vice versa.
- ❑ The time zone and local language.
- ❑ For each user, program folders and other Start menu settings.
- ❑ For each user, all preferences in all user applications unless the applications are older 16 bit applications that use .INI files. Preferences include the “recent files” list on the file menu.
- ❑ All user profiles.
- ❑ User and group security information.

Centralized Configuration Management with the Registry

Windows NT’s Registry keeps everything in one place, eliminating the need for the various types of initialization files of 16 bit Windows. You can view Registry settings by using one tool the Windows NT 4.0 Registry Editor. Windows NT 4.0 provides two versions of the Registry Editor: Regedit.exe, a Single Document Interface(SDI) Explorer-type tool derived from the Windows 95 RegEdit application and Regedt32.exe, a Multiple Document Interface(MDI) version that originated in Windows NT3.1 Regedit.exe offers the advantage of a more flexible Find feature; Regedt32.exe offers advanced administrative features, such as the ability to set security and audit properties of individual keys. The choice of versions is up to you; Windows 95 users are likely to choose Regedit.exe and most users experienced with prior versions of Windows NT opt for Regedt32.exe. This lecture uses Regedt32.exe, except as noted.

Registry Organization

The Registry is arranged in a logical and straight forward way that clearly distinguishes among the following three classes of settings:

- System-wide, for all application and all users (for example, your computer’s microprocessor type)

- System-wide, for one user(for example, your Windows Desktop color Scheme)
- Per application for each user(for example , the last four files you opened in Excel)

Although you continue to use various tools to change Registry values in Windows NT4.0 don't underestimate the value of using just one tool to view any of the settings. Navigation such a large collection of settings and configuration information can be a bit intimidating, but understanding the structure and content of the Registry is very important when administering Windows NT Server4.0.

The Danger of Changing Registry Values

Changing entry values in the Registry can cause your server to be unstartable. Rather than change Registry entry values directly, you should use one of the following tools of change these values:

- Most hardware settings are handled by the hardware recognition process when you start your computer, PCI adapter cards configure themselves automatically during startup. Because Windows NT 4.0 doesn't fully support plug and play for ISA cards, legacy ISA adapter usually use jumpers or semi permanent software configuration methods. Alternatively you can install the pnpisa.sys driver, which lets you statically configure PnP ISA adapters, When you install hardware drivers, driver configuration information is stored in the Registry.
- Many system configuration and application settings are set when you run the appropriate setup program. For this reason, it's important not to relocate applications by simply moving their files. Registry entries aren't updated when you move an application's executable and support files, so you must remove and reinstall most applications change their location.
- The Administrative Tools menu contains User Manager, Disk Administrator, performance Monitor, Backup, Event Viewer, and Windows NT Diagnostics tools. Using these dedicated tools to interact with the Registry is far safer and more intuitive than using the Registry Editor.
- Control panel contains the Accessibility Options , Add/ Remove programs, Console, Date/Time, Devices, Display, Fonts, Internet, Keyboard, Licensing, Modems, Mouse, Multimedia, Network, ODBC,PC Card(PCMCIA),Ports Printers, Regional Settings, SCSI Adapters, Server, Services , Sounds, system, Tape Devices, Telephony and UPS tools, in addition to any tools that are added by application you install . Control panel tools interact with a small part of the Registry for examples, the keyboard tool changes the user's keyboard settings in the Registry.

- Using desktop applications changes the Registry. For example, the list of recent files is generated automatically as you open files in the application.
- Most OLE servers and ActiveX components register themselves when they're installed or run for the first time.

How the Windows NT and Windows 95 Registries Vary

In many important ways, the Windows NT Registry is similar to the Windows 95 Registry. Certainly the concepts are the same. The same tools and applications read from and write to the Registry, and the Regedit.exe Registry Editor of Windows 95 and Windows NT 4.0 is quite similar. However the names for specific collections of information stored within the Registry aren't identical. Nonetheless experience with the Windows 95 Registry is readily transferable to Windows NT Server 4.0s Registry.

15.4 Registry's Organization

Understanding the Registry's requires you to learn another new vocabulary. The Registry is made up of keys, some of which have subkeys. Keys have value entries and groups of keys and their value entries are gathered into a hive. The Registry is a hierarchical database; keys correspond to records in the database.

Registry Hives and Files

The Registry is arranged in a hierarchy quite similar to a folder tree. At the top of the hierarchy are the following five hives:

- HKEY_LOCAL_MACHINE contains system-wide hardware information and configuration details stored in the SAM, Security, Software, and System configuration files.
- HKEY_CLASSES_ROOT Contains OLE and ActiveX information and file associations.
- HKEY_CURRENT_CONFIG contains OLE and ActiveX information stored in the system files that's also kept in HKEY_LOCAL_MACHINE.
- HKEY_CURRENT_USER contains all the settings specific to the current user, which stored in the Ntuser.dat files located in the \Winnt\Profiles\Username folder

- HKEY_USERS contains all the settings for all user, stored in Ntuser.dat files, including the current user and a default user, which is stored in the \Winnt\Profiles\Default User folder.

These five hives are permanent Registry components.

Configuration files associated with Registry hives are stored in the \Winnt\System32\config folder . The four system files that store the contents of the HKEY_LOCAL_MACHINE and HKEY_CURRENT_CONFIG hives. Configuration files use the following extensions:

- Files without an extension hold the current version of the configuration information
- Alt holds a backup copy of the HKEY_LOCAL_MACHINE\System key, which is critical to starting Windows NT. Only the System file has an .alt version
- .LOG contains the transaction log that holds all changes made to the configurations file until the change is made permanent.
- .dat contains user information. Only the new NTuser.dat files use the .dat extension Ntuser .dat in \Winnt\System32\Config replaces the Usernamexxx and Adminxxx files of prior versions of Windows NT .NTuser.dat in \Winnt\System32\Profiles\DefaultUser replaces the previous Userdef file. NTuser.dat.log is the log file for NTuser.dat
- .sav file are created by the text mode part of Windows NT 4.0 Setup and are used in the event that the graphics mode part of Setup fails.

As mentioned earlier, pairing hive and log files ensures that the Registry can't be corrupted. For example, if the power fails as your change to a value entry is written the value might be changed but the date stamp might still contain the old date and time, or the size of the entry might not be correct. The hive and log approach guarantees that these types of error don't happen.

When it's time to write out changes to a hive file, Windows NT inserts a few extra steps . First, Windows NT writes to the log file the new data and instruction for where a special mark at the beginning of the hive file to indicate that it's being changed. The changes are written to the hive file and , on completion, the mark is removed. If a power failure or other serious problem occurs during the process, Windows NT notices when opening the file that the "being changed" mark exists re-creates the changes from the log file, and then removes the mark. This process maintains the consistency of the hive files.

Keys and Subkeys

Just as a folder in a file system can have subfolders a hive has keys and a key can have subkeys. Just as a folder is itself a folder, a subkey is a key, another named collection of information. Each key can have many subkeys, each of which can have even more subkeys, and the hierarchical can be many levels deep.

The HKEY_LOCAL_MACHINE hive has the following four keys, which commonly are also called hives because they have associated configuration files:

- HKEY_LOCAL_MACHINE\SAM is the directory services database (formerly security Account Manager) stored in the SAM files
- HKEY_LOCAL_MACHINE\ Security stores local security information such as specific user permissions, in the security files.
- HKEY_LOCAL_MACHINE\ Software holds configuration information for applications and their components in the software files.
- HKEY_LOCAL_MACHINE\ System stores information that controls startup of Windows NT and loads the necessary device drivers, followed by Windows NT services.

The HKEY_LOCAL_MACHINE\System hive is loaded so early that a simpler process is needed. System .alt is just a copy of the system hive file. The changes to system aren't logged but the "begin changed" mark is still used. After system is written, system .alt is written in the same way. If the power fails while your computer is writing System, Windows NT will notice the "being changed" mark and use System .alt as a safe backup. Of course, the changes being made when the failure occurred are lost.

By Convention, hive, key and subkey names are gathered into full names and separated with backslashes that appear like folder path names. An example is a key called HKEY_LOCAL_MACHINE\hardware \description \system \Multifunction Adapter\0\DiskController\0\FloppyDiskPeripheral. The Floppy Disk Peripheral key is a subkey of the 0 key, which is a subkey of the Disk Controller key, and so on, up to the hive HKEY_LOCAL_MACHINE

Value Entries

To continue the analogy of a file structure further, a key can contain value entries and subkeys just as a folder can contain files and subfolders. Value entries in keys resemble

files in folders. A value entry contains the information to examine or change just as a file contains the data you display or edit. A key can (and often does) support more than one value entry.

A value entry has three components:

- The name of the value entry
- The type of information it contains (numerical or character data, for instance)
- The value of the information (C:\program.exe or 0, for example)

The following sections describe each of these components.

Value Entry Names Microsoft chose reasonably comprehensible names for most value entries; you probably can guess what Current User, InstallDate, LogFile Path and DiskCache Size contain without any need for documentation. Much of the information is added to the Registry when you install programs on the system and application vendors may not choose sensible or easy-to-understand names.

When there's only one value entry in a key, it's possible –but not necessarily wise –that the programmer who added the key left the name unassigned. When omitting a single value name, the Registry Editor shows (default) in place of the missing value entry name. This practice is quite common in the file association entries of the HKEY-_CLASSES_ROOT key. Use the name of the key to understand the information in the value entry.

15.5 Registry Editor

Remember that the Registry Editor isn't found on any of the default menus in Windows NT. Run a special program to begin using the Registry Editor. Here's the command name:

```
% SystemRoot % \ system32\REGEDT32.EXE
```

You can add it to the Start menu for easy access.

Because changing Registry entries can cause damage to your system, however we recommend to review the Registry in Read Only mode until it is familiar with its working. To open in Read Only mode, choose Options Read Only Mode and make sure that a check mark is next to this option (**see figure 15.1**). Using this feature can prevent accidents while allowing us to become familiar with the Registry.

The Registry Editor uses an Explorer like interface to display five cascaded Windows (see figure 15.2). Each of these windows depicts one of the predefined Registry sub trees for Windows NT. Think of these sub trees as the root directories for the five areas of the Registry. The following sections describe the keys.

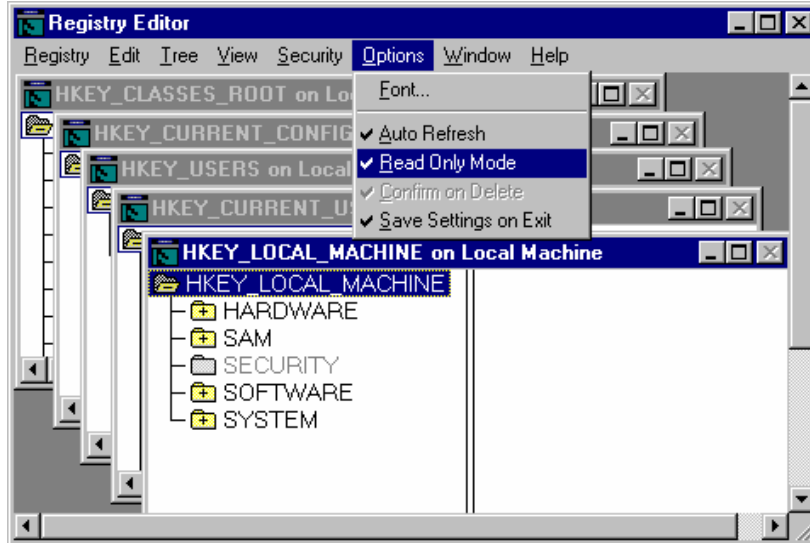


Figure 15.1 The Registry Editor allows users to view the Registry in Read Only Mode

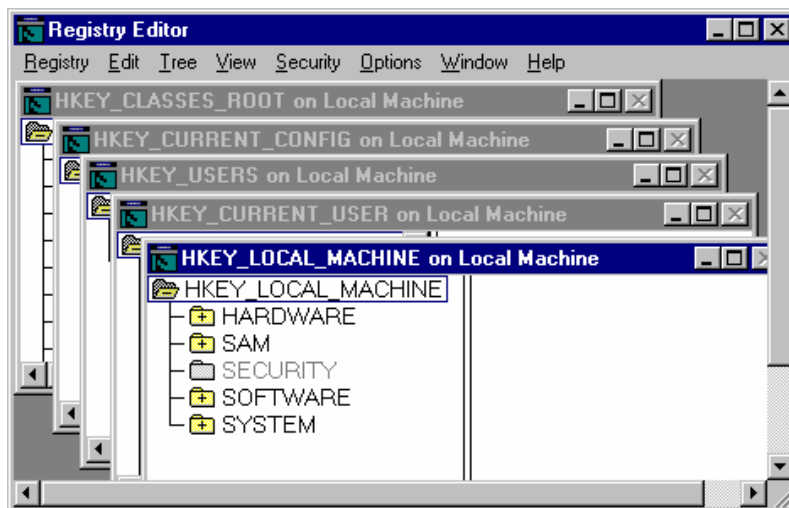


Figure 15.2 The Registry Editor provides a standard interface for viewing and modifying the Registry.

Table 15.1 shows the five predefined sub trees with a brief description

Sub tree	Contents
HKEY_LOCAL_MACHINE	Information about the local computer.
HKEY_USERS	User profiles from users who have logged on to the Local system.
HKEY_CURRENT_USER	The user profile from the user currently logged on.
HKEY_CLASSES_ROOT	Object linking and embedding (OLE) information for the local machine.
HKEY_CURRENT_CONFIG	Information for the configuration used at startup.

Table 15.1 *Predefined subtrees in Windows NT 4.0*

15.6 Important Hives & Keys

It's not possible to list all the keys in your Registry, because keys are added every time you install hardware or software. It's not even possible to list all the keys in the Registry or a typical machine running Windows NT Server and no applications; publication limits on page count preclude such a listing. It's also not useful to provide a list of all Registry keys; once you know roughly where to look, it's quicker to use the Registry Editor to search for the exact key or value entry. Thus, this section doesn't attempt to describe all Registry keys, but it suggests appropriate locations to search for particular classes of Registry entries.

HKEY_LOCAL_MACHINE

The HKEY_LOCAL_MACHINE window holds configuration information about the local computer. This information includes hardware configuration, security data, and details about the software loaded on the machine. The information in the hardware sub key is created dynamically each time the system is started. The program NIDTECT.COM detects the hardware and places the hardware information in the Registry. Because it's re-created at startup, this information isn't saved in the Registry. This also means that any changes to these keys won't be saved.

The Security Account Manager (SAM) holds security information for the local machine. If the machine is a domain controller, SAM holds all the security information for the entire domain. As of Version 4.0, the operating system, by default, will no longer allow you to view or modify the security information. As an administrator, we can change security values to view this information, but it is in a binary format.

The software SUB_KEY contains information about the software loaded on the local machine. This information is usually machine-dependent and doesn't conflict with information in the HKEY_USERS. If there is a conflict with HKEY_USERS, the information in HKEY_USERS will be used.

The system subtree holds information used to start the system. This topic is discussed shortly.

HKEY_USERS

The HKEY_USERS window contains user profiles from users who have logged onto that machine locally.

HKEY_CURRENT_USER

HKEY_CURRENT_USER is a mirror of HKEY_USERS key. When a change is made to one, the change is made to the other. This key holds the user profile of the user who is currently logged on. This window includes information about the current user's desktop settings, Printers, network connections, and software configuration. If there's conflict between an entry in this key and hkey_local_machine, this key takes precedence.

HKEY_CLASSES_ROOT

Normal **Figure 15.3** shows the HKEY_CLASSES_ROOT key, which holds information on object linking and embedding (OLE). This key mirrors the HKEY_LOCAL_MACHINE\Software\Classes subkey. The duplication of information in the Registry is another method that Microsoft used to ensure fault tolerance in Windows NT.

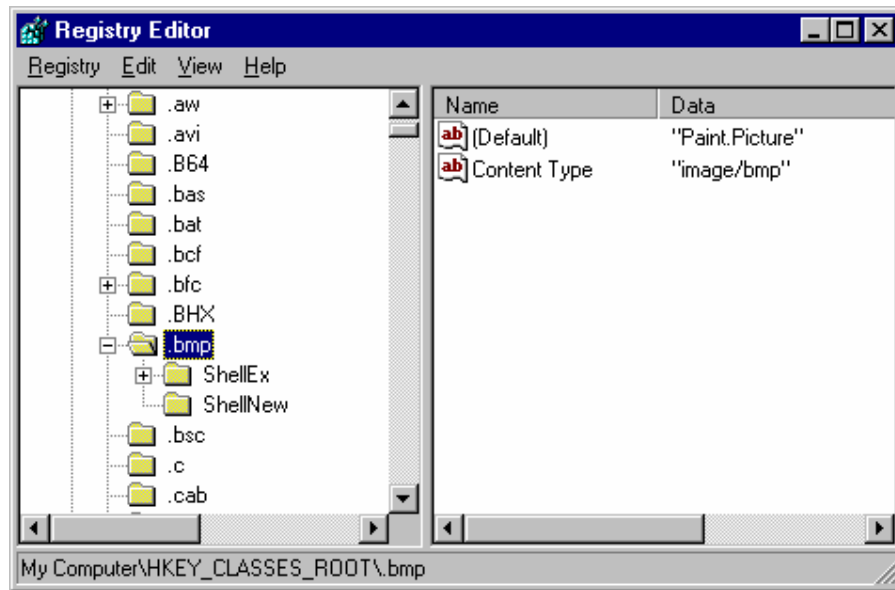


Figure 15.3 The HKEY_CLASSES_ROOT subtree

HKEY_CURRENT_CONFIG

To display information about the current hardware configuration at startup, open HKEY_CURRENT_CONFIG. This key was added in Windows NT 4.0 along with the availability of hardware profiles.

Using the Windows NT Diagnostics Utility

In many of the preceding sections, you read warnings against modifying Registry values. Sometimes, however you need to know the value of a Registry setting, particularly a hardware setting. One way to look at all your hardware related Registry settings at once is to use Winmsd, the Windows NT Diagnostics utility.

As its name implies the Windows NT Diagnostics utility helps you diagnose the behavior of your system by examining various settings at once. To run Winmsd, from the start menu choose programs, Administrative Tools and Windows NT Diagnostics. Winmsd gathers various settings into nine tabbed pages

- *Version.* Operating system information including version number build number serial number and registered owner.
- *System.* Processor and BIOS information
- *Display.* Display type, settings and drivers.

- *Drives.* All fixed removable, and remote drivers arranged by type (floppy, local hard drive, CD-ROM, remote hard drive)
- *Memory.* System memory and paging files.
- *Services.* Running or Stopped status for all system services or devices.
- *Resources.* Ports, mouse, floppy drives. And other system resources.
- *Environment.* Environment Variables.
- *Network.* Domain , workgroup, access level and what user is logged on now . Network settings and statistics are also available.

If you aren't sure which Control Panel tool or Registry key to use to check a Registry value, use Winmsd. Winmsd is quick and you can't accidentally change a value. If you want to keep a record of the settings described in the preceding list, use the printed reports from Winmsd rather than compile a written list from values displayed by the Registry Editor.

Backing Up the Registry

Backing up the Registry often is very important especially before you change anything to try to fix a problem. Some , but not all Registry information is saved on the emergency repair disk. The following are four different ways to back up the Registry:

- From within the Regedt32.exe Registry Editor, choose Save key from the Registry menu and save the key to alternate media such as tape or a drive elsewhere on the network . to restore from this backup, choose Restore Key from the Registry Editor's Registry menu. You must individually select each hive window and save its associated configuration file.

In Regedit.exe choose Export Registry File from the File menu to open the Export Registry File dialog. Select all in the Export Range section to export the contents of the entire Registry to a single to a single text file with a .reg extension .using Regedit.exe to export Registry files is simpler than using Regedt32.exe. Saving the OAKLEAFO Registry resulted in a 5M text files; the first few lines of notepad.

- If you backup to tape by using Windows NT Backup, mark the backup Local Registry check box in the Backup Information dialog, and the Registry is backed up to tape with regular files. To restore from this backup, use Windows NT Restore.
- Use the Regback.exe or Repaire.exe programs included on the Windows NT Resource Kit CD-ROM to backup Registry files, To restore the Registry use Regrest.exe or Repair.exe.

- From another operating system, copy the files in the \Winnt\System 32\Config folder to alternative media. Also copy the user information hives from each \Winnt\profiles\Username folder. To restore the Registry, use the other operating system to copy the backups into those folders again.

15.7 Inspecting Another Computer's Registry

Occasionally you may need to view or change a setting on another computer that's not located with your windows NT Server computer. Fortunately Windows NT server can view and edit Registry entries of other computers running Windows NT or Windows 95. You also can edit the Registry of your server from another Windows 95 or Windows NT machine.

Rather than travel to the computer with the problem or tell the user to make the change with another tool, you might choose to edit the Registry remotely from your computer in at least the following situations:

- The user can't tell you what's wrong and can't search through the Registry to report what settings have been changed.
- The user doesn't have the authority or the skill to change the Registry value, even when you can dictate the necessary keystrokes by telephone.
- The problem is so quick to fix and the remote computer is so far away that it's not worth the time for you to go to the computer.
- The problem is one that makes the computer very hard to see / an example is when the background and foreground text colors have been set to the same color so that no one can read the text in dialogs.

Preparing for Remote Registry Editing Of course, you don't just take over the Registry of any other computer on your network and start changing values.

Editing a Windows 95 Registry from Windows NT If the remote computer is running Windows 95 you must enable user-level security and remote administration from that computer's Control Panel. First, use the network tool's Access Control page and select user-level Access Control. Then use the password tool's Remote Administration page and mark Enable Remote Administration of this Server. Finally, on the Network tool's configuration page mark Add the Microsoft Remote Registry Services.

Editing a Windows NT Registry from Windows 95 If you are trying to manipulate a Windows NT Registry from a Windows 95 machine (perhaps to fine-tune the departmental server from your desktop), you must add the Microsoft Remote Registry services and arrange user-level access; however you don't need to enable remote administration of the Windows 95 machine. No preparation is required on the Windows NT Server machine.

Editing a Windows NT Registry from Windows NT If you need to manipulate the Registry of one Windows NT machine from another (any combination of Windows NT server and Windows NT Workstation), no preparation is required.

Opening the Remote Registry To edit from a Windows NT machine start Regedt32.exe and use the Registry menu's select computer command to open windows with the HKEY_LOCAL_MACHINE and HKEY_USERS keys of the remote computer, alternatively start Regedit.exe and choose Connect Network Registry from the Registry menu. Then inspect or change the keys and value entries to the remote computer. The changes take effect immediately so use extra care especially if the remote computer is in use while you're changing the Registry. You might want to save this task for a time when the remote machine isn't in use or arrange such a time with your users.

From a Windows 95 machine start Registry Editor and choose Connect Network Registry from the Registry menu to open the remote Registry. You can then inspect or change the keys and value entries of the remote computer. Windows 95 client(OAKLEAF1, My computer) connected to the Registries of OAKLEAF0 (Windows NT Server4.0) and OAKLEAF3 (Window 95). The HKEY_DYN_DATA hive appears only for Windows 95 Registries.

15.8 Maintaining Registry Security

You can restrict user's ability to change Registry values in a number of ways. Such restriction should be part of an overall security plan that allows users to access only those administration features they need.

First, don't provide Administrator access to non-administration. You can restrict the access non-administrator users to the Registry. You also should consider deleting Registry Editor and Policy Editor from client computers. On windows NT machine delete Regedit.exe and pldedit.exe from the \Winnt folder, and remove Regedt32.exe from the \win 95 folder . You can administer clients from the server , which usually is in a location with more physical security than the rest of the network or from your client PC.

To control access to individual keys, you can add or remove names from the Access Control List (ACL) for each key. If you care enough about a particular key's value to restrict access you should audit access to the key or audit failed access attempts.

Restricting Access to Registry Keys

The process of restricting access involves several different administrative tools. Only Regedt32.exe provides access to the security properties of keys. Follow these steps to set up auditing and security for one or more key values:

1. In User Manager for Domains, choose policies from the audit menu and make sure that Audit These Events is selected, Select Success or Failure or both for File or Object Access
2. Within Registry Editor Select the key for which you want to restrict access and then choose owner from the Security menu. If you aren't the owner of a key, you can't change permissions for that key. As administrator you can change the owner to yourself, but you can't return ownership unless the original owner gives you full Control permissions on the key.
3. After you confirm or take ownership of the key, choose Permissions from the security menu. The Registry Key Permissions dialog is used to assign permissions to the groups listed. To add another group, click Add; to remove a group click Remove. When the permissions for this key are correct, click OK.

The available permissions are as follows:

- ◆ *Full Control.* Users in the group can view, change, take ownership and change Permissions. Administrators and the system group should have Full Control on every key.
 - ◆ *Special Access.* Users in the group can view and change the key.
 - ◆ *Read.* Users in the group can only read the key.
4. Choose Auditing from the security menu to arrange auditing of key access. The Registry key Auditing dialog appears
 5. Select the types of accesses you want to be logged for each group.

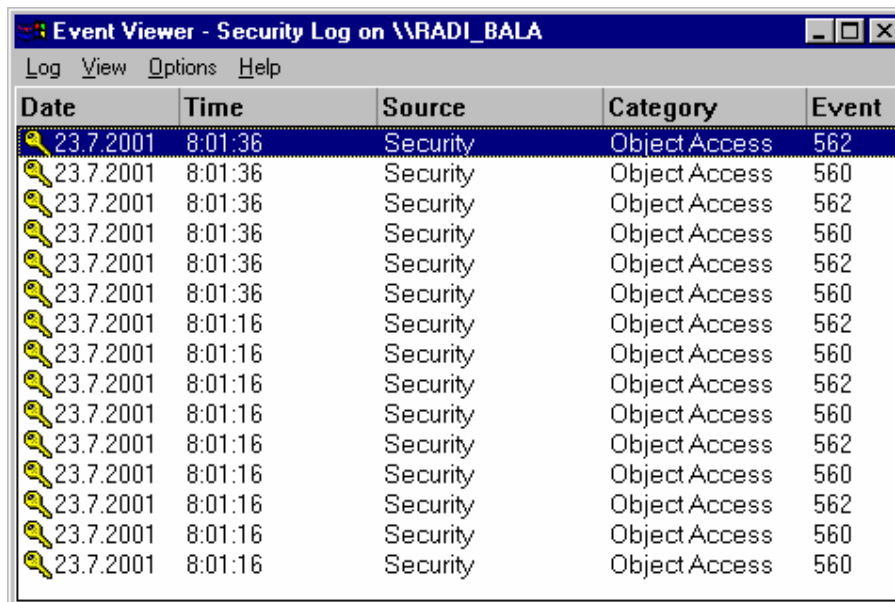
You probably don't want to log successful accesses because there may be a large number of accesses. For example, many keys are updated every time a user runs an application and each update may generate several log entries. Logging failed accesses allows you to discover applications that are no longer working for users, or users who are trying to change keys for which they have no permission.

The types of access audits are as follows:

- ◆ *Query Value.* An attempt to learn the value of the key.
- ◆ *Set Value.* An attempt to change the value of the key.
- ◆ *Create Subkey.* An attempt to make a subkey within the key
- ◆ *Enumerate Subkeys.* An attempt to list the subkeys of this key.
- ◆ *Notify.* Notification events from the key.
- ◆ *Create Link.* An attempt to create a link within a key.
- ◆ *Delete.* An attempt to delete the key.
- ◆ *Write DAC.* An attempt to change the permissions (Discretionary Access Control) on a key.
- ◆ *Read Control.* An attempt to learn the permissions on a key.

Viewing an Audit Log

To view the audit logs, from the start menu choose programs, Administrative Tools and Event Viewer. From the Log menu choose Security to see a list of the logged events. Figure 15.4 shows a sample list. These entries aren't very helpful beyond the user name; double-click one to see details like those in figure 15.5.



Date	Time	Source	Category	Event
23.7.2001	8:01:36	Security	Object Access	562
23.7.2001	8:01:36	Security	Object Access	560
23.7.2001	8:01:36	Security	Object Access	562
23.7.2001	8:01:36	Security	Object Access	560
23.7.2001	8:01:36	Security	Object Access	562
23.7.2001	8:01:36	Security	Object Access	560
23.7.2001	8:01:16	Security	Object Access	562
23.7.2001	8:01:16	Security	Object Access	560
23.7.2001	8:01:16	Security	Object Access	562
23.7.2001	8:01:16	Security	Object Access	560
23.7.2001	8:01:16	Security	Object Access	562
23.7.2001	8:01:16	Security	Object Access	560
23.7.2001	8:01:16	Security	Object Access	562
23.7.2001	8:01:16	Security	Object Access	560
23.7.2001	8:01:16	Security	Object Access	560

Figure 15.4 Event Viewer's Security log listings for Registry events

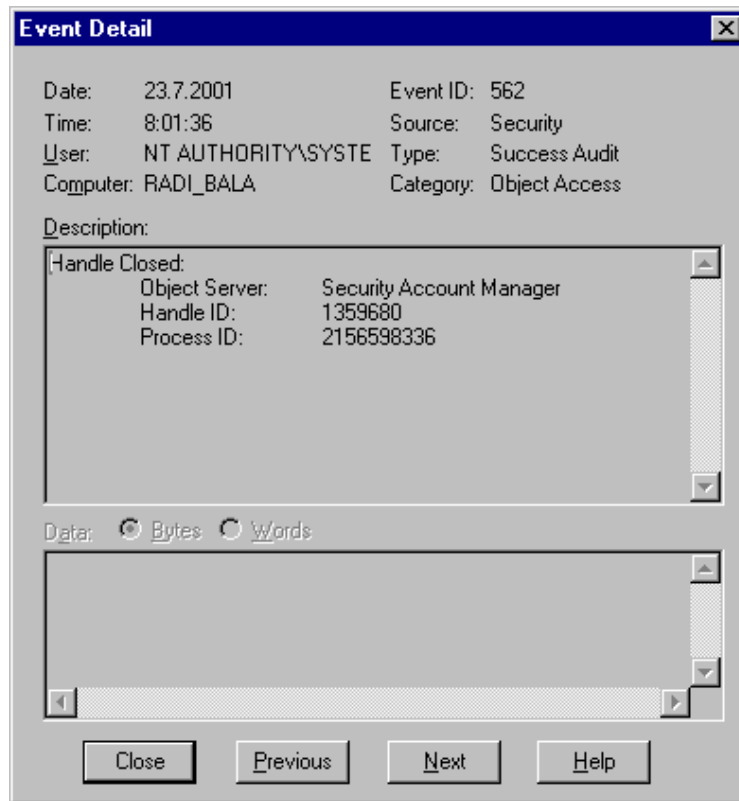


Figure 15.5 Viewing Security Log entry detail for opening a Registry key object

These entries were generated as follows:

- 1 Auditing was turned on in User on for both success and failure for any access to keys under HKEY_USERS \the-SID-of-the-account-in-use-at-the-time\software\Microsoft\Notepad
- 2 In Notepad, the font was changed and then Notepad was closed.
- 3 In the Event Viewer, the menu item Log Security was chosen

Even at a cursory glance, notepad key was changed. Further investigation could narrow down the source of a user's trouble quite easily. For example, if you're logging only failures, you might relax security restrictions so that the operation no longer fails.

15.9 Short Summary

- ☞ The registry is a binary database of the setting that Windows NT and its application need to start and operate.
- ☞ There are five different data types for the Windows NT Registry. Each data type uses a different format.
- ☞ The software sub key contains information about the software loaded on the local machine. This information is usually machine-dependent.
- ☞ A very good way to prevent unauthorized users from modifying the Registry is to Change the Protections on the Registry Editor.
- ☞ REGEDIT doesn't allow you to view or modify security in the Registry.
- ☞ Use REGEDT32 to work with the Registry's security features.

15.10 Brain Storm

1. What is Registry?
2. What are the uses of Registry?
3. What are the Subtrees available in NT?
4. Explain Each Sub trees.
5. How NT Survive in the case of Power Outages and System Crashes?
6. What is the use of Remote Registry Editing?



Lecture 16

Using TCP/IP, WINS and DHCP

Objectives

In this Lecture you will learn the following:

- Understanding the concept of TCP/IP, WINS and DHCP
- Able to configure the TCP/IP , DHCP and WINS clients
- Implementing DHCP and WINS

Coverage Plan

Lecture 16

16.1 Snap Shot

16.2 Role of TCP/IP

16.3 Installing & Configuration TCP/IP
--

16.4 Implementing DHCP

16.5 Implementing WINS

16.6 Short Summary

16.7 Brain Storm

16.1 Snap Shot

This lecture describes TCP/IP and related applications that have gained wide acceptance and use over the last decade. TCP/IP is the network protocol used on the Internet, which by itself makes the topic worthy of study. It's also very useful in private networks, especially as they grow in size. You will learn how to install TCP/IP, DHCP and WINS

16.2 Role of TCP/IP

TCP/IP is a suite of network protocols that describe precisely how information can be transmitted from one computer to one or more additional Computers. TCP/IP is designed to operate in environments where the conditions aren't particularly suitable for this task, and therefore has a strong error-detection and correction capability. Most often, the term TCP/IP denotes not only the protocol suite itself but also a group of compatible applications and utilities that have been created and used to implement and test the protocols.

Members of the Internet community have developed TCP/IP cooperatively by using a proposal and peer-review process involving documents called Request for Comments (RFCs) . A person or group proposes a design and publishes an RFC describing that design. Other members of the community. Some of whom may refine the proposal with their own additions again put forth in an RFC, review it some of these designs are implemented, tested, and refined even further. Eventually, an RFC that describes a set of standard is developed, and manufacturers design produces that confirm to one or more of these RFCs.

This process turns out to be quite effective, over time at discovering and eliminating problems. The RFC process is ongoing and existing RFCs are available for public review. RFCs primarily are-intended for individuals and organizations who design products and services to be used on the Internet. Some RFCs include useful information for Internet users and don't describe standards at all. RFC 1118, The hitch hikers Guide to the Internet, " is an example of this type of RFC.

Why Use TCP/IP?

The use of TCP/IP is growing for many reasons. During the last decade, many organizations implemented LANs in offices and sites throughout their facilities. Eventually, they wanted to connect these LANs into WANs Also growing number of organizations have started to view the WAN as a strategic resource, critical to the success of their efforts. To implement these views they need a protocol that can manage large numbers of systems in a routed, WAN environment. This is precisely what TCP/IP was designed to deliver.

TCP/IP is also the protocol used on the Internet and is therefore useful for individuals and organizations who want to attach directly to the Internet or access it through a service provider. Furthermore, TCP/IP allows a high degree of interoperability between dissimilar systems, such as computers running Windows NT and Unix operating systems. TCP/IP also provides an environment that supports the development of powerful applications having feature-rich programmatic interfaces.

IP Addresses, Host Names, and NetBIOS Names

The central capability provided by TCP/IP as already mentioned, is a transmission facility moving information from point A to point B. The transmission of information must be done in a manner that takes into account the involvement of both computers and humans. The computers must be able to send and receive information accurately and quickly, and their human operators must easily be able to specify what action they want and understand the results.

That computers and humans require different naming schemes for the elements of a network is the source of much of the difficulty surrounding its operation. Computers must have a unique address for each component on the network to accurately send information to just those components for which it was intended. Humans also must be able to specify the computer they want to communicate with, and to name their own computer system so that they can describe it to other humans, especially if they're sharing information on the network. But the kind of name appropriate for computer use is much different from what's suitable for humans.

This leads to one central problem that TCP/IP must solve—name and address resolution. Three types of name are designed for humans, and two addresses are designed primarily for computers and their operating systems and applications. Matching a name with its corresponding address is more difficult than it might at first appear. A Windows NT network using TCP/IP uses the following name types:

- *Machine address*, also called hardware address. In Ethernet networking, the machine address is a guaranteed unique address that's "hard wired" or manufactured into a computer network product, such as a network adapter for a personal computer. Ethernet network adapters use a Media Access Control (MAC) address, which includes a portion that's specific to a particular manufacturer so that two different manufacturers never create the same address. Within their private address space, each manufacturer must make sure that it never creates two devices with the same address. This is usually done by including a ROM chip or similar element with a unique identifier that becomes part of the address. MAC addresses are expressed as 12 hexadecimal digits (for example, 00 04 AC 26 5E 8E) often written with a space

between each two digits for human readability. Other hardware networks, such as ATM and token-ring networks, use different schemes to assign machine addresses.

- *IP address* used by operating systems and networking software on TCP/IP networks. If you create a private network, you must make sure that no two devices have the same IP address. If you want to attach to the Internet, you must request part of the address space from Internet NIC (Network Information Center) for you're Organization to use and then manage that portion so that no two components use the same address. .

IP addresses are written in a form known as dotted decimal notation. For example. 123.45.67.89 is a valid IP address. Each part is called an octet and can range from 1 to 254 (0 and 255 are generally reserved for special purposes). This address must be unique for each device on a given network. It's composed of two parts, the network ID and the host ID. The network ID, the first two octets, must be the same for all devices on particular network segment or subnetwork and different from all other subnetworks. The host ID, the last two octets, must be unique within a particular network ID.

- *Host name*, the 'human -compatible' name for a computer or device on a TCP/IP network. A host name is also called an FQDN (Fully Qualified Domain Name), or simply a domain name when specified in full. A host name for a server might be data server, and its FODN might be dataserver. Company com. Applications using host names are generally case sensitive. You can use this name instead of the IP address when entering many commands for TCP/IP specific applications and utilities FQDN's aren't used when entering windows based Microsoft networking commands, such as NET use or NET VIEW, which require a NetBIOS name.
- *Domain name*, another name for the host name. The last part of this hierarchical name (company com for example) is referred to as a first level (or top-level) name and is used to uniquely identify your organization to the Internet community. Often a request for a domain name in an application or operating system utility refers only to the first-level name net the FQDN
- NetBIOS name used for Microsoft networking commands, such as NET USE and automatically used on your behalf when performing networking functions with windows-based graphical utilities, such as File Manager or the Windows 95 Network Neighbourhood. A NetBIOS name can be 15 characters (for example, DATASERVER) Applications using NetBIOS names aren't generally case sensitive.

Resolving Names and Addresses

During the execution of a network command, the application or operating system must eventually discover the machine address of the devices involved. Because users almost never enter the machine address into an application, you must use some means of resolving the host name, NetBIOS name or IP address to machine address. Various mechanisms for this purpose have been developed and are discussed in this section.

Separate mechanisms exist for each type of name, and sometimes more than one process may occur. For example, an application that knows the host name may first resolve the name to an IP address and then to a machine address. The mechanisms for resolving each type are presented in the following list and are discussed in more detail in the next section. Some of these mechanisms are based on standards as defined in RFCs or other standards documents, and others are Microsoft specific methods. IP addresses are resolved to machine addresses by using the following methods.

- Address Resolution Protocol (ARP), defined in RFC 826
- A search of the corresponding ARP cache in the computer's memory

Host names are resolved to IP addresses by using the methods in the following list. If the computer is configured to use all methods, the methods are tried in the following order;

- HOSTS file
- Domain Name Service (DNS)
- Windows Internet Name Service (WINS)
- A local broadcast
- LMHOSTS file

NetBIOS names are resolved to IP addresses by using the methods presented in the following list. If the computer is configured to use all methods, the methods are tried in the following order.

- A NetBIOS name cache in the computer's memory
- WINS
- A local broadcast
- LMHOSTS file
- HOSTS file
- DNS

Name-resolution mechanisms for host names and NetBIOS names are similar but are carried out in a different order. The mechanisms used can vary, depending on how the computer is configured.

16.3 Installing & Configuring TCP/IP

This section shows you how to install TCP/IP on a Windows NT server, including all the TCP/IP options Windows NT Server 4.0 offers. Microsoft's implementation includes various client based utilities for the TCP/IP suite, such as Finger, lpr, rcp, rexec, rsh, Telenet, and tftp.

Both client and server support is provided for FTP (File Transfer Protocol). FTP allows as Windows NT server or workstation to interact with UNIX workstations and other platforms supporting TCP/IP. Notably missing from the connectivity utilities is support for NFS, the Network File System that makes files on UNIX servers accessible to PCs running Windows, but NFS is available from third-party software companies for Windows NT servers. A number of diagnostic utilities also are offered; arp, hostname, ipconfig, lpq, nbtstat, netstat, ping, route, and tracert. An SNMP agent, implemented as a Windows NT services, enables use of a remote network management console, such as Sun Net Manager or HP Open View, with your server.

You can install TCP/IP during the original setup of Windows NT server or add it later from Control Panel's Network tool. In this section, you add TCP/IP to an existing Windows NT server 4.0 installation. Adding it during the initial setup is an almost identical process, so the following steps should still be helpful. You simply follow these instructions when you get to the network portion of Setup.

To install TCP/IP and related services, follow these steps:

1. From the Start menu, choose setting and control Panel.
2. Double-click the Network icon to open the Network property sheet. On the Protocols page, click the Add button to open the select Network Protocol dialog. Select the TCP/IP Protocol item in the network Protocol list and click OK.
3. A dialog requesting the full path to the distribution files appears. Enter the location by using the drive letter of your local CD-ROM drive with the Windows NT Server 4.0 distribution CD_ROM inserted or the Universal Naming Convention name of a shared network resource with the distribution files. (the UNC name option is feasible only if you have another network transport protocol already installed and operational). Click continue to copy the required files to you server.

The Microsoft TCP/IP properties sheet, shown in the below figure appears on completion of the copy process. (if the Microsoft TCP/IP Properties sheet doesn't appear, double click the newly added TCP/IP Protocol item in the Network Protocols list.)

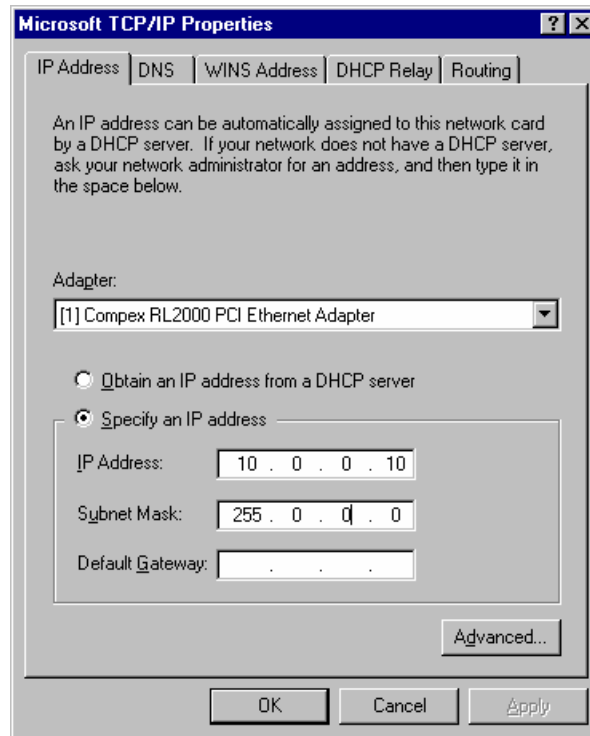


Figure shows the default IP Address page of the Microsoft TCP/IP Properties sheet

4. Use the IP Address page to set the IP address for this server. You can setup an IP address in two ways:

- Obtain the address via DHCP “Dynamic Host Configuration Protocol” select the Obtain an IP Address from a DHCP server option. The DHCP server provides all the needed information. DHCP is primarily intended for assigning IP addresses to client workstations. .
- Statically assign an IP address, which is the most common method for servers . You must obtain a unique dotted decimal IP address from a central authority on your network. This person also can provide the subnet mask. The default Gateway is the main router used to forward packets to and from other networks, most likely the Internet.

5. Use the Advanced button to configure additional TCP/IP settings for additional adapters. A computer with more than one network adapter is called a multihomed computer.
6. The DNS page lets you to enter one or more DNS servers for name resolution . Windows NT has already set your host name to the machine name that you picked for your computer. Enter a valid domain name for instance, radiant.com, if your network has DNS server installed. The domain suffix search order lets you search multiple domain name spaces.
7. If WINS is enabled on your network, you can enter the IP addresses of primary and secondary WINS servers on the WINS Address page. If a DNS server is available on the network, marking the Enable DNS for Windows Resolution check box causes the TCP/IP protocol to check the DNS server first for name resolution.
8. The routing page has one check box, Enable IP Forwarding; marking it enables static routing and if installed dynamic routing on this computer.
9. Click OK to commit the changes and close the property sheet . You must reboot the machine so that all the changes take effect.

In addition to the base protocol, You may want to install some of the TCP/IP –specific services. These services include TCP/IP printing (which lets users print to TCP/IP printers) , simple TCP/IP and SNMP, (SNMP, OR Simple Network Management Protocol, lets you configure and monitor network devices.) These services are most useful when you print to existing UNIX host printers. The simple TCP/IP services install an FTP server and a Telnet server. The SNMP services allow this computer to become the source of SNMP events or the destination of SNMP messages.

16.4 Implementing DHCP

Setting up a DHCP server requires defining a scope, configuring client reservations, configuring DHCP clients, testing clients, viewing and managing DHCP client leases, and maintaining your DHCP database. A DHCP scope is a pool of available IP addresses and (optionally) additional addressing information for various shared devices or services. As a DHCP client computer connects to the network a unique IP address is assigned and with the addresses of other shared resources (for example, servers) can be transmitted to the client computer.

The *IP address* is said to be *leased* to the client computer because it can be returned to the pool of available addresses and used by another client later. You can define global

options that apply to all scopes defined on a DHCP server. You can also define options that apply to only one scope.

You need several pieces of information before you can complete the configuration of DHCP. You must answer the following questions before configuring DHCP:

- Will all the computers on your network be DHCP clients? If not, you must be sure to exclude the addresses from the pool of available addresses. In general servers, routers, and other similar devices should be configured with static IP addresses.
- What information, in addition to the IP address, do you want to configure automatically? A default gateway? WINS server? DNS server?
- What options can you configure for all clients on the network? What options do all clients on a particular subnet share? Are any options unique for specific clients?
- How many DHCP servers do you need? If your network consists of multiple physical subnets connected by routers, your routers must act as BOOT Relay Agents as specified in RFC 154, or you must put a DHCP server on each subnet with DHCP clients. If your router doesn't support RFC 1542 (many older routers don't), you may be able to upgrade it to add such support without having to replace the router.
- What range of addresses and other information should you include in the scope defined on each DHCP server? Should you define multiple scopes for any servers? Remember DHCP servers don't share information with other DHCP servers. Each must have its own set of addressees to offer to the clients it will service.

The Advantages of DHCP

DHCP offers several advantages over the manual configuration of TCP/IP addresses:

- Network users aren't required to enter an IP addresses, subnet mask, or any other addressing information. Therefore they're much less likely to enter a random address or copy an address from a colleague's computer (following the reasoning that if a colleague's addresses works an identical configuration also will work on their own computer).
- The process of manually entering an IP addresses, subnet mask, and other configuration information is prone to error, even with educated user population that cooperates full with the process. There are too many numbers and settings to expect a large group of users to set them without error. When users change computers or locations, the settings must be redone.

- A fair amount of administrative overhead is associated with managing the list of valid IP addresses, even with DNS. This process is also inherently difficult to divide among several individuals unless each is knowledgeable about the technology and cooperates fully with one another.
- DHCP lets users configure their own computers without having to contact an administrator to get valid IP addresses. This eliminates errors, delays and frustration.
- When users move their computer to a new location or travel with a laptop containing a PCMCIA Ethernet adapter or similar device, they automatically receive a valid address for the new location when they start their computers.

How DHCP Leases IP Addresses to Clients

Before you install the DHCP server an over view of the DHCP lease address process may help you administer the process more effectively. Following is a basic description of how DHCP leases IP addresses:

1. A client computer starts and initializes an unconfigured version of TCP/IP. Then it broadcasts request for an IP address. The request contains the computer's hardware address and computer names so that DHCP server knows who sent the request.
2. All DHCP servers with an available lease that's valid for the client send a response by using a broadcast message (because the client doesn't have an IP address yet). The message includes the client's hardware address, the IP address being offered the subnet mask, the duration of the lease, and the IP address of the server making the offer. The server must reserve the address in case the offer is accepted.
3. The client accepts the first offer receives. It broadcasts its acceptance to all DHCP servers with a message including the IP address of the server whose offer was accepted. Other servers release the temporary reservation on their offered addresses.
4. The server with the selected address sends an acknowledgment message with the IP address, subnet mask and possibly other information defined in the scope as described in the next section. The client receives the acknowledgment and initializes a full version of TCP/IP. The client now can communicate with other hosts on the LAN or WAN.

DHCP Server Configuration

You use the DHCP Manager utility to configure a DHCP server. To open DHCP Manager, from the Start menu choose programs, Administrative Tools , and DHCP Manager You can start, stop pause and continue, the DHCP service as you can all services-by using the Services icon in the Control panel or by using the Windows NT Server Manager. Make sure that the service, formally named Microsoft DHCP Server, is started.

The rest of this section describes the procedures you use to define a DHCP scope, set various option and configure and test DHCP client workstations. You also can reserve certain manually assigned addresses (for example, for servers and routers) so that they're excluded from the pool of available addresses managed by the DHCP service.

Creating a DHCP Scope

To create a DHCP scope follow these steps:

1. Start DHCP Manager as explained in the preceding section.
2. From the Scope menu choose Create to display the Scope dialog.
3. Enter the range of IP addresses to include in this scope it's usually a good idea to include the full list of addresses used on this network or subnet and then explicitly exclude those addresses managed by a DNS or other DHCP server. You may also want to set aside a range of addresses for servers, routers, or other network devices so that you can establish addressing conventions that make it easier to identify shared devices by their Ip addresses. For example. Within a given scope, you might set aside host Ids from 1 to 20 for servers and 250 to 254 for routers and hubs, even if you don't need them all at this time.
4. Enter a subnet mask. If you aren't subnetting, the class of your IP address determines the mask For example, 255.255.255.0 would be used for class C addresses suitable for small networks with fewer than 255 hosts .
5. Enter a name for the pool and optionally include a descriptive comment. Set the lease duration based on the volatility of your host population. For example, if you have a very stable network, set a long duration If you have a small range of addresses that an ever changing group of traveling laptop users must share, set a short duration,.

6. Click OK A dialog informs your that the scope has been defined but not activated You can activate it now or wait and activate it later by highlighting the scope and choosing Activate from the Scope menu.

A DHCP server would be implemented on each subnet . Each DHCP server can back up the other with a range of addresses from the other scope.

Configuring Global Options To set options that are provided to all clients from all scopes as they receive an IP address lease follow these steps :

1. Start DHCP Manager, if necessary.
2. From the DHCP options menu choose Global to open the DHCP Options: Global dialog.
3. In the DHCP Options Global dialog, select an option from the Unused options list Click Add to move it to the Active Options list.
4. Select the Option in the Active Options list, and click Value to expand the DCHP Options: Global dialog
5. Click Edit Array to open the IP Address Array Editor dialog
6. Enter the IP addresses of elements that correspond to the option you selected in step4 which is shown in the General Information section. Click Add.
7. Repeat steps 3 through 6 for each option you want to apply globally to all scopes.
8. Use the arrow buttons in the IP Address Array dialog to arrange the entries from the top down order in which you would like them to be used (not all options are consulted in this order, depending on the nature of the options used). Click OK twice to return to the DHCP Manager window.

Configuring DHCP Clients You configure Windows clients to use DHCP as follows:

- Windows NT Workstation 4.0 Launch Control Panel's Network tool, click the Protocols tab, double-click the TCP/IP Protocol item and select the Obtain an Address from a DHCP Server option on the IP Address page of the Microsoft TCP/IP Properties sheet.
- Windows 95 Launch Control Panel's Network tool, Which opens with the Configuration page of the Network property sheet active. Double-click the TCP/IP → Network Adapter item and select the Obtain an IP Address Automatically option on the IP Address page of the TCP/IP Properties sheet.
- Windows for Workgroups Run Network Setup and mark the Enable Automatic DHCP Configuration check box in the TCP/IP configuration dialog.

All other settings can be received from the DHCP server if they're defined in the scope used by this client. Any entries made for other parameters (the default gateway, for example) take precedence over values received from the DHCP server.

Testing DHCP Clients In this section, you use the Ipconfig diagnostic command-line utility to report the status of your current network configuration. Use Ipconfig to view the IP address you've leased from a DHCP server and other information passed to your computer from the defined scope. To verify the operation of DHCP view your current address release it, and then renew a lease. This operation is only for testing or other diagnostic and troubleshooting use.

To test the operation of a DHCP client, follow these steps:

1. Start the client computer and log on the network.
2. Open a command prompt. Type the following command:

```
ipconfig /all
```

This command displays a full listing of your IP address and all options that were defined globally, for your scope, or for your individual client workstation.

3. If you've defined options for DNS servers, WINS servers, a default gateway, and so on, try using the ping command with their addresses. This command "bounces" a test packet off the other machine and returns it to your computer to test basic network connectivity. For example, by using the address of a WINS server defined in the examples used for the figures you would enter.

```
ping 182 . 111 . 200.3
```

You should receive a series of replies with the time it took to make the trip to the remote host and back. Ping other devices configured for your scope or globally on your network.

4. Enter the following command to release your IP address:

```
ipconfig / release
```

5. Re-enter the command

```
ipconfig /all
```

6. You, no longer have an IP address and can't communicate with other hosts on the network. Now enter the following:

```
ipconfig /renew
```

This command renews your lease, probably with the same address (unless another host happened to lease it while it wasn't being used). Check the information you received from the DHCP server by using the /all option with ipconfig again.

Maintaining the DHCP Database The DHCP database, DHCP.Mdb (like the WINS database. WINS.mdb) uses a modified version of the jet (Access) database file structure. Jet databases increase in size when updated, because replaced or deleted records are marked for deletion, not physically deleted. At periodic intervals you should compact the DHCP database by using a command-line utility provided for that purpose. Jetpack.exe reclaims wasted space in the database left by the update process. For large networks, this maintenance should be performed approximately once a week; for smaller networks, once a month is appropriate. You must stop the Microsoft DHCP Server service before you can perform this operation, so this operation is best done during off-peak times.

To use Jetpack exe to compact the DHCP database , follow these steps:

1. Use Control panel's Services tool or Windows NT server Manager to stop the service. YOU can also use the following command: `net stop dhcpserver`
2. Open a command prompt and change to the `\systemroot\System32\DHCP` directory
Make a backup copy of the database, just in case it's needed:
`copy dhcp.mdb dhcp.bak`
3. Use Jetpack to compact the DHCP database, creating a new temporary file that replaces the existing database: `jetpack dhcp.mdb temp.mdb`
4. Delete the existing database (remember you have a backup copy): `del dhcp.mdb`
5. Rename the compacted temporary database as in-use database: `ren temp.mdb dhcp.mdb`
6. Restart the service: `net start dhcpserver`

16.5 Implementing WINS

WINS is Microsoft's implementation of a NetBIOS name server (NBNS) . It's implemented as a Windows NT Server service. With the WINS Manager administrative utility program and appropriate client software WINS register the NetBIOS Name used by client and server computers as they start. When a Microsoft networking command, such as NET USE, initiates a networking operation WINS handles the subsequent need to resolve a NetBIOS name to complete the command WINS also resolve TCP/IP host names after the local HSOTS file is checked and the DNS server is consulted.

The Advantages of WINS

WINS dramatically reduces the amount of broadcast traffic on the network. Because name resolution with WINS is handled by direct communication between WINS clients broadcast name registration requests and name query requests are therefore minimized.

You also don't need to configure all clients to use WINS –you can operate a mixed environment. WINS resolves names from clients across routers and can therefore support multiple subnets.

If a WINS server isn't available the design of the system still enables when a WINS server is down. WINS servers can replicate the names in their databases to other WINS servers so that a single dynamic names database is represented and managed across the enterprise network.

How WINS Registers and Resolves NetBIOS Names

To use WINS you must configure a WINS server and start the service whose formal name is listed in the Services dialog simply as Window Internet name Service. The next section covers the steps involved in configuring a WINS server.

After you set up one or more WINS servers and WINS enabled clients, the process of registering and resolving names involves a number of distinct processes that are carried out in a natural order.

Before you configure a WINS server, it's helpful to understand the configuration of the client and how WINS name resolution happens. The steps involved in WINS name registration follow:

1. A WINS client is configured with the address of the primary and an optional secondary WINS server. The WINS server names can be specified on the client or received with an IP address as one of the optional DHCP parameters passed from a DHCP server. As the client starts, it sends its NetBIOS name directly to the WINS server in a name registration request.
2. If the WINS server is available and the name isn't already registered to another client, the registration is successful, and a message is returned to the client with a positive registration and the amount of time for which the name is registered, known as the *Time to Live* (TTL)
3. If a duplicate name is found, the server sends a name challenge to the currently registered client. If the client responds and affirms that it's using the name the new registration is denied by sending a message to the requesting client. If the currently client doesn't respond to three queries the name is released and registered to the new client.
4. If the ARP client can't find the primary WINS server after three attempts, an attempt is made to find the secondary WINS server (if the client has been

configured for a secondary WINS server). If the secondary WINS server also can't be found with three ARP requests, the client resorts to a standard broadcast to register its name with the local subnet.

By default, a WINS client uses the h-node implementations of NetBIOS over TCP/IP. The process involved in WINS name resolution follows:

1. When actions in the Windows interface enter or implicitly specify a command a name resolution is required. The Net BIOS Name cache is checked first to see whether the NetBIOS name mapping to an IP address is available.
2. If the mapping isn't in the NetBIOS name cache, a name resolution query is sent directly to the primary WINS server. If no response is returned, the request is sent three times.
3. If the primary WINS server doesn't respond, the secondary WINS server (if configured) is tried as many as three times. If either the primary or secondary WINS server receives the request, it looks up the name in its database and sends the IP address back to the client or replies with a Requested name does not exist message if it's not listed in the database.
4. If the name can't be resolved by a WINS server (because the server is unavailable or the name isn't in the database) a, h-node name resolution query is broadcast as many as three times.
5. If the name still isn't resolved, the LMHOSTS (if configured) and HOSTS files are searched.
6. If the name isn't in LMHOSTS or HOSTS, the DNS (if configured) is consulted.

Wins configuration

A single WINS server can resolve for an entire WAN because the request are sent as directed datagrams and can be routed. A secondary WINS server provides redundancy and fault tolerance. Additional WINS servers can be provided based on the number of client requests received and performance considerations in large network environments. A rough rule of thumb is that a

typical WINS servers can handle as many as 1200 name registrations and 700 name queries per minute. A pair of WINS servers should be able to handle as many as 8,000 WINS clients under typical network conditions. If you implement servers with two or more processors, a pair of WINS servers should handle more than 12,000 clients.

WINS servers don't need to be domain controllers as well they should be configured with a static IP address subnet mask, default gateway address, and other TCP/IP options. Using DHCP assigned options is possible but not recommended.

Entering Basic Configuration Information.

To configure the basic operation of your WINS server, follow these steps

1. From the Start menu, choose Programs Administrative Tools, and WINS Manager to open WINS Manager's Window.
2. From the Server menu choose Configuration to open the WINS Server Configuration dialog.
3. The WINS Server Configuration section contains settings for time periods that control the basic behavior of the server: the length of a name is registered and how often a client must register its name. For most installations the default values are appropriate. Click OK after making any adjustments.
4. From the Option menu choose preference dialog
5. The settings you make in the preference dialog control the address display and refresh rate of the statistics display. Make any changes you want to configure the display to suit your needs. For example you might want to select an Address Display option that shows both the NetBIOS name and the IP address. Click OK.

Configuring WINS Clients

You can configure WINS clients by simply entering the address on a primary WINS Server and optionally a secondary WINS server, into the client's configuration. You can do this manually by using control Panel's network tool or automatically by using DHCP using DHCP you can manually configure can manually configure individual clients and those settings take precedence over the DHCP settings. If you use DHCP to configure WINS addresses you must also configure clients with option 046 WINS/NBT Node type otherwise. WINS won't work with DHCP. Set option 046 WINS/NBT Node type to Ox8(h-node or hybrid)

Maintaining the WINS Database

You use the Jetpack command line utility used to compact the DHCP database to compact the WINS database as well. It's recommended that you compact the database if it grows larger than 30M. You must stop the Microsoft WINS Server before performing this operation so backup is best done during off-peak times.

To use Jetpack to compact the WINS database, follow these steps:

1. Use Control Panel's Services tool or the Windows NT Server Manager to stop the WINS service. You can also use the following command: `net stop wins`
2. Open a command prompt and change to the `\systemroot \SYSTEM32\WINS` directory. Make a backup copy of the database just in case it's needed as follows:
`copy wins.mdb wins.bak`
3. Use Jetpack to compact the DHCP database creating a new temporary file that replaces the existing database: `jetpack wins.mdb temp.mdb`
4. Delete the existing database (remember, you have a backup copy:) `del wins.mdb`
5. Rename the compacted temporary database as the in-use database: `ren temp.mdb wins.mdb`
6. Restart the service with the following command: `non start wins`

16.5 Short Summary

- TCP/IP is a suite of network protocols that describe precisely how information can be transmitted from one computer or more additional computers.
- TCP/IP is designed to operate in environments where the conditions aren't particularly suitable for this task
- DNS is an IP address resolution method frequently used on UNIX systems.
- DNS servers are implemented and can then be consulted to resolve names not listed in the local HOSTS file
- DHCP is a protocol that allows IP addresses to be assigned automatically from a pool of available IP address centrally stored and managed on one or more servers.
- The IP address pool managed by a DHCP server must be entirely owned by that server.
- Microsoft has designed a server-based service, an administration utility and client software that implement the DHCP protocol
- WINS is a NetBIOS Name Server(NBNS) implemented as a NT service.
- WINS is a dynamic name service that tracks network names as users start and stop client workstations.

16.6 Brain Storm

1. Give a notes on IP addresses, Host and Domain Names.
2. Explain the role of TCP/IP?
3. How to resolve the Names and Addresses?
4. Write the steps to install and configure TCP/IP for NT server?
5. How you can implement Dynamic Host Configuration Protocol?
6. Write the pros of DHCP?
7. How DHCP Leases IP Addresses to Clients?
8. Write the steps to configure WINS Server?
9. How WINS Registers and Resolves NetBIOS Names?
10. How you can maintain DHCP Database?



Lecture 17

Working with Domains

Objectives

In this Lecture you will learn the following:

- Able to implement the domains and trusts between domains
- Knowing the Domain Models

Coverage Plan

Lecture 17

17.1 Snap Shot

17.2 Win NT- Domain Models

17.3 Domain Architecture & Security

17.4 Implementing Domains and Trusts between
Domains

17.5 Short Summary

17.6 Brain Storm

17.1 Snap Shot

One primary strength of Windows NT Server is its domain facility for providing access to and control of network resources wherever needed. Windows NT Server provides file, print and application services to a variety of clients in various environments. Access to these services doesn't need to be limited by geography or network bandwidth. More importantly, control and administration of these services can be distributed to fit either a centralized or decentralized support mode. Windows NT Server uses domains to provide unfettered access to and administration of resources in large, distributed networks.

This lecture emphasizes the value domains bring to your Windows NT network environment and how to take best advantage of domains in LANs and WANs.

17.2 Domain Models

When you begin the design of your Windows NT network environment, you can choose from four standard domain models to use for your organization, as well as hybrids of the standard models. Before you decide on a specific model, carefully analyze how the model matches the business methods of your organization. When you've installed a large network infrastructure, changing your domain design requires a major effort. An impossibly large number of ways exists to organize your network into domains, but Microsoft claims that the possible groupings all fall into one of four basic categories. There are five domain models

The Single-Domain Model

In a single-domain environment there's only one primary domain controller, which is shared by all machines in the domain (There should be a backup controller, but it doesn't count because it just mirrors the primary controller). This lone controller keeps all user account and access information for all machines within the domain. The Single Domain model doesn't use trust relationships because there's only one domain.

The Single Domain model is best suited for an organization with centralized administration, a homogeneous user population and less than 5,000 users. Beyond 5,000 users it makes sense to start breaking up domains into resource and account domains.




Figure 17.1 *Single Domain Model*

The single-Domain model is the simplest domain-based model. It requires the least overhead of any of the domain model, but it is best suited for specific circumstances.

Advantages:

- Good for smaller network
- Administration of any server from any machine in the domain
- High security
- Splitting of single domain into multiple domain is easier

Disadvantages:

- Loss of granularity cause a loss of network performance due to browsing
- Every browse query must go to the PDC for resolution
- Users see long lists when they browse the domain.

The Master Domain Model

The master domain model, the network has a separate “master” domain that contains all users, group and security data. The “slave” domains, which can refer all logon requests to the master domain’s controller. In this model, user accounts are defined as global accounts by the master domain; the master domain administrator can group accounts in global as needed. In turn each of the sub domains can choose to which user accounts a global groups to grant access.

Advantages:

This approach maps well to the common situations of having central MIS or IS group that owns network access. The IS group can create and manage network wide user

account and individual departments can still control which of those accounts can use their resources. The sub-domains can still create local accounts that allow intra-domain access to their data.

Disadvantages:

Single domain system that controls logon and group data is the establishment of that domain as checkpoint. Because all logon requests have to go to the master domain, network traffic between segment increases. The master domain model is not well suited to use over WAN links or networks where there's slow connection between the subnets containing the master and slave domains.

The Multiple Master Domains Model

Multiple master domains solve the choke point problem that plagues the master-domain model by having more than one master domain. All the master domains trust one another and each subdomain trusts all the master domains.

The Complete Trust Model

The Complete trust model is simple to explain, but can be difficult to administer. Each domain has individual trust relationship with every domain it want to trust. The number of trust relationships required to do this can grow very quickly – a network with 4 domains requires 12 trust relationships to enable full trust between all domains, but a network with 5 domains needs 20, 6 domain network requires 30, and a 10 domain network requires 90! The maximum number of relations for a set of N domains is expressed by calculating $n \times (n-1)$. Trust relationships are one-way so this maximum reflects full bi-directional trust between all domains, which not every network needs.

Advantages

This model allows exact control over trust relationship, even in very large domains. This model also does away with centralized control, because each domain can decide which other domain to trust.

Disadvantages

This model cause a lot of additional work for each domain' administrators. Operationally, the lack of centralization may cause problems if each group is responsible for maintaining its own relationships, because relations often need to be two-way.

Hybrid Domain Models

In addition to the proceeding four models, you can combine two or more of the standard models to create a hybrid model that best fits your organization. For example you might like the concept of the Multiple Master domain model, but you don't want the Master Account Domains to trust each other. Instead you want to create a number of Single Master Domains, each with its own resource domains. You also want to give a subset of users in your organization, such as network administrators, access to all domains everywhere as in the Complete Trust.

You can build a new Master Account Domain, called Shared for users requiring cross-resources access. In this case all your resource domains also are trusting to the cross-resource domain. Users requiring cross- resource access log on to the Shared domain and are granted access to as many resource domain as needed.

Many other options exist for domain models. The only caveat is that one-way trust relationships aren't transitive-that is, if Domain A resources are trusting Domain B resources are trusting Domain C users, there's no relationship between Domain A resources and Domain C users.

17.3 Domain Architecture and Security

In the simplest terms, a Windows NT domain is a logical grouping of servers, workstation users, groups and printers within a physical network. The architecture of Windows NT domains is quite complex, involving security issues, different types of user groups, and truest relationships. For this introductory discussion, you can consider a domain to be a group of resources (Windows NT objects) that are bound by a common membership into a single administrative unit. The purpose of domains is to permit distributed network management by segmenting the resources of large network into sets of manageable sizes.

If you have used Windows 95 or Windows for Workgroups in a network environment, you're familiar with the concept of small groupings of users who need access to common resources, such as files and printers. These resources reside either on members PCs or on a central peer-to-peer file server. Domains are extensions of workgroups, designed to support larger sets of users who may be located at geographically distant sites.

Windows NT Server authenticates users when they log on to a server in a domain. Authentication is the process by which users gain access to the resources for which they have rights. Unlike earlier network operating systems, such as Novell NetWare 3.X,

users log on once to the domain, a rather than log on to each server they want to access. Likewise, administrators can assign right to a Domain User or Domain Group, and these rights are applied universally to any domain resource users want to access.

Understanding Windows NT Security Identifiers

Beyond the concept of grouping users, machine, and server resource into logical domains, windows NT provides a mechanism to secure a resource's right to the domain. That is, when a server or client joins a domain, or when a user or user group is created within a domain, windows NT server provides a mechanism to guarantee that the new resource is uniquely associated with the domain. From a security standpoint, this mechanism also assures that a user or machine that isn't properly identified to the domain can't access resources in the domain. Windows NT uses a SID in a addition to a name, to identify each domain resource, group, or user. The SID is generated at the time of creation of the resource and is unique. You can't duplicate a SID, because it is based on a variety of CPU information at the moment of its creation.

When you establish a new domain by installing Windows NT server as the Primary Domain Controller (which is the default for your first Windows NT Server installation) the domain receives its own SID. As you add servers, clients, groups, and users to the domain, their unique SIDs include a reference to the original domain SID. If you move a client PC from one domain to another, the client's SID references the wrong domain and can't use the resources of the new domain.

The advantage of a secure logical grouping of resources become apparent when you must manage large groups of users and machines. Network security is paramount in Windows NT Server's domain architecture. Fortunately, windows NT Server's administrative tools, such as User Manager for Domains and Server Manager, shield you from most of the complexity of SIDs. One of Microsoft's guiding principles in the design of Windows NT is that features to make network administration more convenient must not compromise network security in any way.

Understanding the Roles of Domain Controllers

Windows NT Server supports the following two types of domain controllers both of which are network servers.

- Primary Domain Controller Only one PDC server can exist within a domain.
- Backup Domain Controller You can have any number of BDCs in a domain. The number of BDCs in a domain depends on the number of users in a domain and the structure of your network.

Setting up a Primary Domain Controller. The PDC is by default the first server you install in a new Windows NT network. When you install a new Windows NT server in an existing network and choose to create a new domain, the Setup program scans your network for a domain controller with the domain name you specified. If no PDC for the domain name is found, the server becomes the PDC.

The PDC's role is critical to Windows NT networking. The PDC is the keeper of all user, group, and machine account information for the domain. This information is stored in the SAM (security accounts manager) database, which resides in files in the `\winnt\system32\config` folder on the PDC's fixed disk. As of Windows NT 4.0, the official name for the SAM is directory database. The PDC is responsible for maintaining the master version of the domain SAM. When you change user passwords, add or remove user and group accounts, or add or remove machines in the domain, the PDC's SAM records these changes.

Changes Backup Domain Controllers

If the PDC is the only domain controller in your domain, it's responsible for maintaining the domain SAM and performing routine domain tasks, such as authenticating user logon requests and maintaining user accounts. If you have a large number of users in a domain, routine domain management tasks can occupy a substantial percentage of server resources, slowing normal file, printer, and application server operations.

If servers and clients are connected to a remote PDC by a low speed network connection, such as Switched-56 or ISDN lines, authentication by the PDC can become very slow. To avoid problems that arise from overload of remote PDCs, Windows NT supports the concept of Backup Domain Controllers (BDCs). A BDC is a Windows NT server installed into a domain after the PDC is installed. The BDC's function is to offload some of the routine domain related tasks from the PDC and provide redundancy in case the PDC becomes unavailable. The BDC contains a copy of the domain SAM, which is replicated (copied) from the PDC on a periodic basis.

The BDC has a local copy of the SAM, so the BDC also can authenticate users, if you have a geographically dispersed network, you assign a BDC to serve remote users. If users want to change their passwords or an administrator wants to create a new user group, the BDC handles the change, and then passes the new password or user account to the PDC during the synchronization process. If the PDC is shut down or otherwise unavailable, you aren't allowed to make changes, such as user passwords or new accounts, to the domain SAM. Even if the PDC is not available, you can continue to authenticate existing users and access resources on the BDC and other servers in the domain.

Promoting a Backup Domain Controller to a Primary Domain Controller

When the PDC fails or is unavailable for an extended period of time, or if you want to change PDC machines, you can promote any BDC in the domain to the PDC. The promoted PDC takes over the role of maintaining the master copy of the SAM database, and the old PDC becomes a BDC just as any other. To promote a BDC to a PDC, follow these steps.

1. Launch Server Manager by choosing Programs and then Administrative Tools from the Start menu.
2. If the PDC is operational, highlight the BDC you want to promote and choose Synchronize with Primary Domain Controller from the Computer menu to assure that the BDC's SAM copy and PDC's SAM are identical. Click OK to acknowledge both messages you receive during the synchronization process.
3. Highlights the BDC you want to promote and choose Promote to Primary Domain Controller from the Computer menu.
4. If the current PDC is operating, you are prompted to acknowledge that you are promoting a BDC in place of the PDC. Click Yes if you want to promote the BDC. If the current PDC is not operating (the icon is grayed in Server Manager), you are warned that the PDC can't be contacted, and that promoting the BDC will conflict with the original PDC, if it becomes operational. Click Yes if you are willing to accept this condition.
5. You see a series of messages indicating that the NetLogon service is stopping and restarting on each of the two servers. When the promotion completes, the previous BDC indicates that it is now a Primary Domain controller, and the prior PDC becomes a Backup Domain Controller.

Using Non-Domain Server. You can choose to add servers to your network that aren't domain controllers. To do so, select the Server option during installation of Windows NT server 4.0 on the machine. Non-domain servers use Windows NT workstation security. You install non-domain servers primarily for the following reasons:

- The server is used for heavy duty applications, such as a database or messaging server, and you want to devote all server's resources to application serving, not domain administrative duties.
- It is certain that the server will be installed in a domain other than the domain in which it's originally set up.

Understanding the Domain Synchronization Process

To spread the load of authentication across BDC's that may reside anywhere on your network, including WAN's with slow or unreliable links, Windows NT performs a periodic domain synchronization process. By default, every five minutes a backup domain controller sets up a connection with the PDC, sends all SAM changes that originated on the BDC during the interval, and receives changed information from the PDC's SAM.

The PDC keeps track of the revision level of the SAM for each BDC in the domain; thus the PDC sends only the incremental changes needed to keep the SAMs synchronized. If the BDC loses communication with the PDC for an extended period, the PDC performs a full synchronization by copying the BDC's changes to the PDC's sending a complete copy of the PDC's SAM to the BDC. This process ensures that the BDC's database is complete and up to date.

Manually synchronizing domain SAMs. If you make a large number of changes to the domain SAM or make changes that must take effect immediately (such as unlocking a user's account), you can perform a full or partial manual synchronization with Server Manager. Server Manager's Computer menu offers the following synchronization options.

- Synchronize the entire domain For a PDC selected in Server Manager's Compute list, choose Synchronize the Entire Domain to cause the PDC to contact each BDC and, depending on the BDC's revision level, send the appropriate records to ensure that the BDC has a complete copy of the current SAM database. A full synchronization is called a push process.
- Synchronize a single with the PDC For a selected BDC, choose Synchronize with Primary Domain Controller to force that BDC to contact the PDC and request all changes since the last synchronization (a pull process). You receive the message when using Server Manager to synchronize a BDC with a PDC. This pull process guarantees only that a selected BDC is synchronized; it doesn't affect any other BDC in the domain.

Changing Automatic Synchronization Intervals: if you're concerned that a large number of SAM changes might affect network performance, you can change Registry values on each BDC to control the replication frequency and the percentage of system resources assigned to perform synchronization.

You add the following new values to the `\HKEY_LOCAL_MACHINE\SYSTEM\Current ControlSet\Services\NetLogon\Parameters` key a BDC or PDC to control how much data is sent from the BDC to the PDC and how often SAM replication from the PDC is performed.

- *ReplicationGovernor* is a BDC value that governs how much data and for how long a BDC communicates with the PDC while synchronizing. The value uses the `REG_DWORD` data type and is expressed as a percentage. If the `ReplicationGovernor` value is missing, the default is 100 (percent). A value of 100 specifies that the buffer the BDC allocates for SAM changes is 128L and the call to the PDC for synchronization occupies 100 percent of the available time to complete the transfer. With a 128K buffer, the BDC can hold about 4,000 SAM changes. If you want to limit the size of the SAM buffers (and , thus, the time a synchronization requires to complete,) you can adjust the `ReplicationGovernor` value for each BDC. For example, on a slow link where few domain changes occur, you might be able to tune the `ReplicationGovernor` value down to 50. if the value is too low, however, the BDC may never be able to complete the synchronization process.
- *Pulse* is a PDC value that governs how often the PDC automatically sends SAM changes to the BDCs. `Pulse` is a value of type `REG_DWORD` for the PDC only and defaults to 300 seconds (five seconds). You can change the value to suit the frequency of SAM updates that occur on your network. Any changes made to the SAM between pulse intervals are sent to BDCs with out of date revision levels.
- *PulseConcurrency* is a PDC value of type `REG_DWORD` that specifies how many BDCs PDC can contact concurrently for updates. This entry defaults to 20 concurrent updates sent to the BDCs. (if you have less than 20 BDCs you don't need to add this value.) the maximum value is 500. the higher the number, the less time it takes to synchronize the entire domain, because the PDC sends more concurrent synchronization. High numbers put a much greater load on the PDC and prevent it from performing other tasks during the synchronization process, such as user authentication or user account maintenance.

Adding Backup Domain Controllers to a New Domain

Even in a small, self-contained domain, you should have at least one BDC to accommodate user logons in case the PDC fails. If you're building a large domain, you must determine the number of BDCs needed to accommodate all your users. Microsoft recommends a maximum of 2,000 users per BDC, but your network architecture is more likely to determine this number. For instance, if you have a branch

office that connects to the rest of your network with a slow link it makes sense to place a BDC in the remote office for local authentication as well as file and printer sharing. Fortunately, it's a relatively simple matter to add BDC as required to handle the authentication load.

Adding Windows NT Clients to the Domain

After you install the PDC and one or more initial BDCs, you add the network clients to your domain. When a Windows client is added to the domain, Windows NT server creates a domain machine account for the client. The machine account is a unique SID assigned to the client name to identify the client to the PDC and BDC (s) so that users logging on to the client can access domain resources. Before a client running Windows NT Workstation 4.0 is installed to a domain, it's part of its own workgroup. The following figure shows the NTWS1 client installed to the workgroup WORKGROUP. NTWS1 maintain its own local SAM database of user and group accounts, which you can view User Manager by choosing Programs and Administrative Tools from the Start menu.

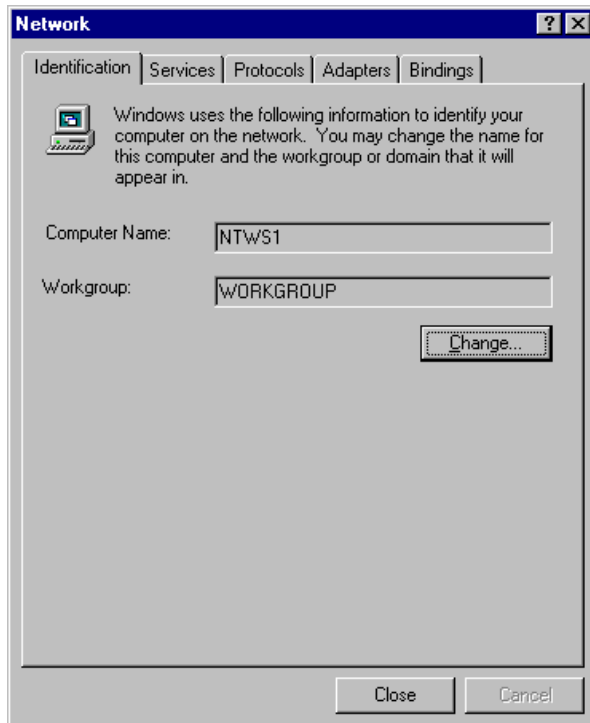


Figure 17.1 A Windows NT 4.0 client installed as a member of a workgroup

When you join a PC client running Windows NT Workstation to the domain, the client's local SAM database is "hooked" into the domain SAM database of the BDC or PDC. The domain administrators global group automatically is made a member of the client's Administrators local group so that Domain Administrator member have access to the client's resources. When you log on to the client, the local SAM authenticates the logon name and password for access to the client's resources. Next the logon name and password to the BDC or PDC for domain authentication and access to domain resources. This process, called workstation security permits logging on to client without the need to join the domain.

Moving a Domain controller to Another Domain. Moving a domain controller from one domain to another isn't a step to be taken lightly. If you want to rename a domain, all domain controllers must be renamed. You can change the domain name of a PDC or BDC to a new, unique domain name (as described in the next section), but changing the domain name doesn't change the domain SID. Thus, you can't move a PDC or BDC to a new domain simply by changing its domain name; a truly new domain requires a new SID, which requires reinstallation of Windows NT server 4.0. after you reinstall Windows NT server 4.0, you also must reinstall as upgrades other applications running on the server, such as Exchange server or SQL Server.

If you must a server that's now functioning as the PDC from Domain A to existing DomainB, follow these general steps.

1. Verify that the Domain Administratory groups has full access to all drives, folders, and files on the PDC to be move. The Domain Administrators group exists in all domains.
2. Promote a BDC in DomainA to PDC. The PDC to be moved becomes a BDC.
3. Perform an independent full backup of the PDC in DomainA, including the Registry. You can restore the full backup in case you reconsider moving the PDC and want to reinstall it as a BDC in DomainA.
4. Perform a new Windows NT Server 4.0 installation on the old PDC into DomainB as a BDC, if the domain exists, or as a PDC if the domain is new. Run Winnt32.exe from the Windows NT server 4.0 CD-ROM and install a new version of Windows NT server 4.0 to a different folder, such as \winnt 41, in an existing partition with sufficient free space for the installation.
5. Start the newly installed BDC. The BDC synchronizes with DomainB's PDC on startup; after synchronization is complete, you can promote the new BDC to PDC.

6. Reinstall as upgrades any server-based applications on the new BDC or PDC.
7. Remove and re-create all file, folder, and printer shares by using the local copy of Explorer on the server, or a local or remote copy of server manager.

When you move a Primary domain Controller from one domain to another, all reference to user and group accounts that existed in DomainA (such as file and folder permission for NTFS volumes) are lost. File and folder permissions for the server in the new domain appear as Account Unknown. Groups and users in the new domain are identified with the new domains SID as a prefix to the group and user SID.

Renaming a Domain. As noted in the preceding section, the original DomainA SID remains as the DomainB SID when you rename a PDC from DomainA to DomainB. Renaming a domain requires you to change the domain names of all devices installed in the renamed domain. Changing the domain name of every device, including BDCs and clients, may be a huge task if you've installed Windows NT to many servers and clients across a dispersed network. If DomainA participates in trust relationships with other domains, renaming the domain breaks the trust relationships, and you must re-create them.

If you must change a domain name, follow these steps.

1. First in Server Manager, highlight the PDC and select synchronize Entire Domain from the Computer menu.
2. Beginning with the PDC, start Control Panel's Network tool. On the identification page of the Network property sheet, click the change button to open the identification change dialog.
3. Type the new domain name in the Domain Name text box and click OK. The warning appears. Click yes to confirm the name change. You are required to restart your server after the change is confirmed.
4. Repeat steps 1 through 3 to rename the rest of your domain controllers.

Rename all your Windows NT Workstation machine accounts for the new domain name.

17.4 Implementing Domains and Trusts Between Domains

The objective of Windows NT server's domains and trusts is to provide users with a single logon to authenticate them to Windows NT server resources, no matter where

these resources are physically and logically located. An understanding of the authentication process is necessary to take maximum advantage of Windows NT Server 4.0's domain and trust architecture. Familiarity with the authentication process also is necessary to optimize the logical topology of your network to balance logon speed with the performance of applications running on your windows NT 4.0 servers.

Distributing Authentication Services

One principal role a PDC and BDC (s) provide in a domain is authentication of Windows NT client and server machines and all users, including users of Windows 3.1 + and Windows 95 clients. Domain controllers verify that a particular user or machine running Windows NT has valid access to the domain.

The authentication process is provided by the NetLogon service, which runs on Windows NT servers and clients that run Windows NT Workstation. The NetLogon service provides a secure channel for message associated with authentication and domain synchronization. A secure channel is a connection between the NetLogon services on two machines that's established and maintained internally by windows NT s security subsystem, using its own set of user names and passwords. Administrators have no access to this connection.

Authenticating Windows NT Clients When a client running Windows NT workstation is powered up, it goes through a series of steps to authenticate to the domain in which it is installed. The authentication process depends, in part, on the network protocol in use. TCP/IP the recommended protocol for new Windows NT networks, a takes the following steps to authenticate a Windows NT client when the Windows Internet Name Service is installed:

1. The Windows NT client installed as a p-node or h-node type WINS client stores its domain name locally and queries the WINS server database first for all domain name entries of 16th Byte type <1c> for the client domain.
2. The WINS server responds with a list of up to 25 domain controllers in its IC listing for the client's domain. This list is dynamic and represents the last 25 domain controllers for the client's domain that have registered with the WINS server, regardless of the domain controller's location on the network.
3. The client sends a NetLogon mailslot packet to each domain controller in the WINS list. The packet, in effect, asks, "Can you authenticate me to the domain?" For h-node type clients, after the mailslot packet is sent to each 1C type server, a broadcast packet is sent out on the local segment asking the same question.

4. All domain controllers receiving the request packet respond, if able, and the first affirmative response to arrive at the workstation handles the authentication request. A domain controller that resides on the same physical network segment as the requesting client is likely respond the fastest.

After the machine is authenticated to the domain by the fastest responding domain controller, users that log on to that machine use the responding domain controller's domain SAM database to verify their identities by means of user names and passwords.

The process a client uses to choose an authentication domain controller is based on the fact that the controller responding most quickly to the request provides the authentication service to the client. There's no simple means to specify particular domain controllers to provide authentication to a set of workstations. If you have domain controllers across a slow link, there's a finite (but very low) probability that one of those domain controllers might respond first to client requests and provide authentication services across that slow link. This situation is undesirable if the link is very slow. The goal is to prevent unwanted responses from remote domain controllers. In a routed TCP/IP environment where WINS provides the list of candidate domain. You can restrict candidate domain controllers by creating a static mapping in WINS of a domain name type with a list of only those domain controllers that you want to respond for authentication services.

In the case of a network that's broadcast based, such as NetBEUI or NWLink, you can control where the broadcasts go by using bridge filters and limiting broadcast forwarding between network segments. Any domain controllers that reside on the same physical network segment as a client, however, always receive the broadcast request for authentication.

Understanding Memory Requirements for and Loads on Domain Controllers. Microsoft recommends the following amounts of RAM for PDCs running Windows NT server 4.0

- For up to 8,000 user accounts, at least 32 M of RAM is needed.
- From 8,000 to 20,000 users, 64 M of RAM is the minimum.
- Beyond 20,000 users and up to about 40,000 users, 148M of RAM is required.

Memory requirements increase if most of your BDCs also must perform other tasks, such as file and printer sharing or application serving. Using the PDC and BDC (s) for ordinary network services is common in smaller networks serving 100 or fewer users. As the number of users increase, bursts of authentication and domain synchronization activity temporarily affect the performance of other services provided by PDCs and


BDCs. Users may notice a significant performance drop in client/server applications, such as database front ends to SQL Server or the Exchange client for Exchange Server, during these bursts. Fortunately, it's relatively easy to add a new BDC to the domain to distribute the server load and to back up the existing BDC, in the event you must promote the existing BDC to a PDC. If you have multiple BDCs consider demoting the BDC running server based application to a plain non-domain server.

Understanding Trust Relationships

If you support a large organization with multiple departments or divisions that want to manage their own resources, use Windows NT server's trust relationships. A trust relationship connects two or more domains and lets users in one domain access resources in another domain. A single logon provides user access to domains with appropriate trust relationship.

You often see domain diagrams depicting one or two-way trust relationships (with arrows pointing in one or two directions, respectively between domains (as shown in the below figure)

One – Way Trust



Two – Way Trust




Figure 17.2 A diagrammatic view of one and two-way trusts

Domains are described as either trusted by or trusting another domain. Resources in the trusting domain are accessed by user accounts residing in the trusted domain. The following figure shows the one-way trust relationship between a trusted and a trusting domains.

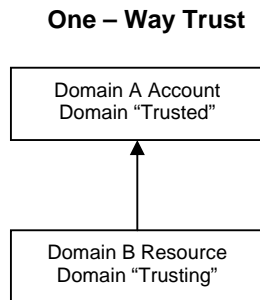


Figure 17.3 A one-way relationship between trusted and trusting domain

A trust lets an administrator in the resource domain assign, for example file permissions to users or groups in the account domain from a file utility such as Explorer. A trust provides only the connectivity between two domains. You must explicitly grant access to resources in the trusting domain for users in the trusted domain, in the same manner that you grant access for users in their own domain.

Trusts allow you to distribute management of resources between multiple domains. For example your IS department may want to manage creation of all accounts and certain centralized server resources in the IS-MASTER domain, while providing users and administrators in the Accounting department the ability to manage resources in their own domain, called Accounting. A one-way trust relationship with Accounting as the trusting domain and IS-MASTER as the trusted domain accomplishes this objects. The accounting users log on to the IS-MASTER domain but are permitted to access and manage resources in Accounting by means of the trust relationship.

Establishing Trusts Establishing trusts is a relatively simple process. By using the User Manager for Domains utility, you can create one and two way trusts between domains. As administrator, you need to have access to a Windows NT Server domain controller in each domain that's part of the trust in order to establish the trust relationship. Follow these steps to create a one-way trust relationship between the account and resource domains.

1. From the account domain, launch User Manager for Domains by choosing Programs and Administrative Tools from the Start menu.
2. Choose Trust Relationships from the policies menu.
3. In the Trust Relationships dialog click the Add button next to the Trusting Domains list to open the Add Trusting Domain dialog.

4. Enter the name of your resource domain in the Trusting Domain text box, then provide and confirm a password in the Initial Password and Confirm Password text boxes. This password must be supplied when you configure the resource domain's trust relationship to the account domain. Click OK to confirm the trusting domain
5. Repeat steps 1 and 2 on a Windows NT server in the resource domain. Then from the Trust Relationships dialog, click the Add button next to the Trusted domain list. Enter the name of the account domain here, and enter the password you specified in step 4 in the Password text box to establish the trust. Click OK to confirm the trust.

Windows NT Server then tries to contact the trusted domain to establish the trust. If you're creating a two-way trust, repeat steps 1 through 5. However, starting on the resource domain, define the account domain as the trusting domain. Then from the resource domain, add the account domain to the resource domain's trusted domain list.

When the trust is established, you have access to users and groups in the account domain from any domain utility you run in the resource domain, such as Explorer, User Manager for Domains or Print Manager.

17.5 Short Summary

- Domain-based security Windows NT organizes servers, printers, workstation, users and groups into one or more domains
- You can designate Windows NT serves as Primary or Backup domain Controllers and promote a backup to a Primary domain Controller when necessary
- Moving a domain controller to another domain or remaining a domain involves a substantial effort
- Windows NT Server offers you to choice of single, single or multiple master, complete trust and hybrid domain trust topologies
- Choosing the right domain architecture for your network depends on the number of servers and clients, administration methods and the type of network connections between domains.
- A two-way trust between two domains, indicated by a two-headed arrow, allows users and groups in either domain to access resources in the other domain. In this case , the two domains are both trusted and trusting.

17.6 Brain Storm

1. Give a notes on PDC.
2. List the types of domain models
3. Write the difference between PDC and BDC?
4. Explain briefly about the trust relationships?
5. Write the pros and cons of one way trust?
6. How you can distribute authentication Services in the domains
7. Write the steps to move a domain controller to another domains
8. Give a notes on the Domain Synchronization Process
9. Explain briefly the domain Architecture and Security.



Lecture 18

Administering User Accounts

Objectives

In this Lecture you will learn the following:

- Understanding the concept of User Account
- Able to manage the User Accounts & Properties
- Able to create the New User account
- Understand the Domain Account Policy

Coverage Plan

Lecture 18

18.1 Snap Shot

18.2 Working with User Manager for Domains

18.3 Managing User Account & User Account
Properties

18.4 Administrating the Domain Account Policy

18.5 Short Summary

18.6 Brain Storm

18.1 Snap Shot

The fundamental purpose of a network operating system is to create a productive environment of users while maintaining a high level of security. This is also the primary goal of all network administrator. This lecture describes about the User Manager for Domain and how to create, copy & modify the users. The last section shows how to administer the Domain Account Policy.

18.2 Working with User Manager for Domains

You can employ User Manager for Domains to manage accounts within any domain that users have administrative access to Individual users have administrative access if they're members of any of the three Windows NT user groups in below Table.

User Groups with Permission to Administer User Accounts and Groups

User Group	Description
Administrators	A local group whose members can perform all user and group management functions.
Domain Admins	A global group that, in most cases, is a member of the Administrators local group. Users in the Domain Admins group are automatically given local Administrator privileges.
Account Operators	A restricted account whose members can manage most properties of user accounts and groups. Members of this group can't manage the following Windows NT Server groups: Administrators, Domain Admins, Account Operators, Server Operators, Print Operators, and Backup Operators. Members of this group also can't manage the account of domain administrators and can't alter domain security policies.

The following sections describe how to take full advantage of the User Manager for Domains application.

Starting User Manager for Domains

You can start User Manager for Domains (called *User Manager* from here on for brevity) from taskbar or from the command line.

To start User Manager from the taskbar's Start menu, choose Programs, Administrative Tools, and User Manager for Domains to open User Manager's window. By default, information for the domain where your user account is defined appears in the window (figure 18.1).

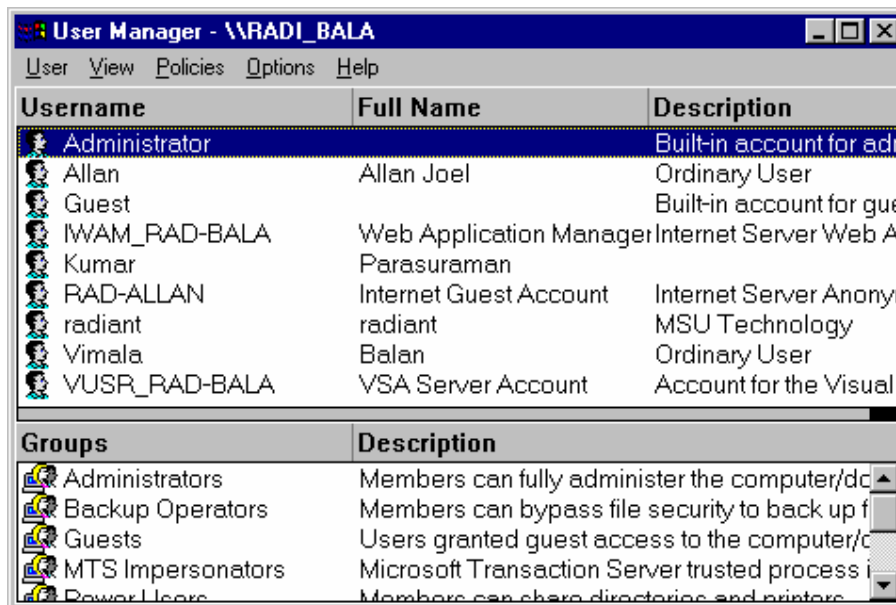


Figure 18.1 Viewing User information for the default domain in the User Manager's main window

Starting Multiple Instances of User Manager

Unlike applications that are limited to a single instance, such as Windows NT Explorer, User Manager allows multiple simultaneous instances. Multiple instances of User Manager is a valuable time-saving feature for administrators of large networks or multiple domains.

The most effective method for running multiple instances of User Manager is to create program icons for each domain or computer that you admin. Each program icon contains the name of the domain or computer as the command-line argument to the program command line. By creating multiple instances of User Manager in this manner, you can admin each domain simply by double-clicking the program icon.

The easiest method for creating multiple copies of User Manager with assigned domains is to follow these steps:

1. From the Start menu, choose Programs and then Explorer to open an instance of the Windows NT Explorer.
2. Open the Administrative Tools folder by moving to \Winnt\profiles\All Users\Start Menu\Programs, and then double-clicking the Administrative Tools folder.
3. Select the User Manager shortcut icon
4. From Explorer's Edit menu, choose Copy.
5. Select the destination folder. To create an additional Start menu item in the Administrative Tools folder, don't change folders.
6. From the Edit menu, choose Paste to add a second User Manager shortcut to the selected folder.
7. Right-click the new User Manager shortcut icon. From the pop-up menu, choose Properties to open the User Manager for Domains Properties sheet.
8. Edit the command-line entry in the Target text box by adding a space and the domain name or the computer name to the end of the command line, as in
%SystemRoot%\system32\usrmgr.exe domain name.

Selecting a New Domain with User Manager

If you choose to manage multiple domains of computers from the same instance of User Manager, changing domains is an easy task. To select a new domain or server, follow these steps:

1. From the User menu, choose Select Domain to open the Select Domain dialog
2. Select a new domain from the Select Domain list by clicking the domain name item. Optionally, you can type a domain or server name in the Domain text box. If you enter a server name, remember to follow the UNC naming convention and precede the name with two backslashes.

3. Click OK to display the user and group information for the selected domain or server.

18.3 Managing User Account & User Account Properties

Every user of a Windows NT Server network must have a user account, which consists of all the information that defines a user to the Windows NT network. The user account defines the resources on Windows NT computers and domains that the user can access.

A user account consists of the typical user name and password, as well as how, when, and where a user can attach to the network; what resources the user can access; and what security rights the user has for the accessible resources. The user account also defines the local and global groups of which the user is a member.

The following sections describe the built-in user accounts, how to add new accounts, and how to modify account properties to take full advantage of Windows NT Server 4.0's support for networked users.

Managing the Built-in User Accounts

When you install Windows NT; two built-in accounts - Administrator and Guest - are established when a domain is created. Unlike named user accounts, the Administrator and Guest accounts can't be deleted. These two accounts are installed on the primary domain controller.

The Administrator Account. The Administrator account is set up by default to allow the installer to manage and configure the Windows NT Server 4.0 software immediately after installation. The user who manages the domain's overall configuration uses the Administrator account. The Administrator account has more control over the domain and its servers than any other user account on the Windows NT network.

During installation of the primary domain controller, the Windows NT Server 4.0 Setup program prompts for the password of the built-in Administrator account. Remember and protect this password. If you forget or lose the Administrator password, the Administrator account is unusable.

The Administrator account is added as a member of the following built-in user groups:

- ◆ Administrators (local group)
- ◆ Domain Admins (global group)
- ◆ Domain Users (global group)

The Administrator account can't be removed from these built-in groups. Some Administrator capabilities include managing security policies; establishing domain trust relationships; creating, modifying, and deleting user accounts; creating shared directories and printers; modifying operating system software; and installing and updating devices and partition disks. This is only a small of the capabilities available to an administrative account with full Administrator privileges.

The Guest Account The Guest account is at the opposite end of the permissions spectrum from the Administrator account. The Guest account is provided for occasional or one-time users. The built-in Guest account is a part of the Domain Guests built-in group and inherits a very limited set of permissions from the group.

Although the Guest account can't be removed from the system, it is disabled by default during installation. This means that you must explicitly enable the account for it to be used. In practice, this account usually is enabled only if some network resources must be accessible by individuals without formal accounts that enable file and other resource sharing. For example, persons who don't need to access server files might be allowed to use the Guest account to use a shared printer.

Adding New User Accounts

Your network isn't useful without users, and it's equally unsecure and unproductive if you use only the two built-in accounts. This means that new user accounts must be added for each network user, with the possible exception of Guest users. Following are the two methods for adding a new user account:

- You can create a new user account from scratch

- You can copy a new user account from an existing user account and make the appropriate changes to specify information specific to the new user.

Creating a New User Account with User Manager. Add a new user account by choosing New User from User Manager's User menu to open the New User dialog. Fill in the dialog's text boxes, mark the appropriate security check boxes, and click the Add button to create the account. Figure 18.2 shows the New User dialog with text box entries, before security options are selected

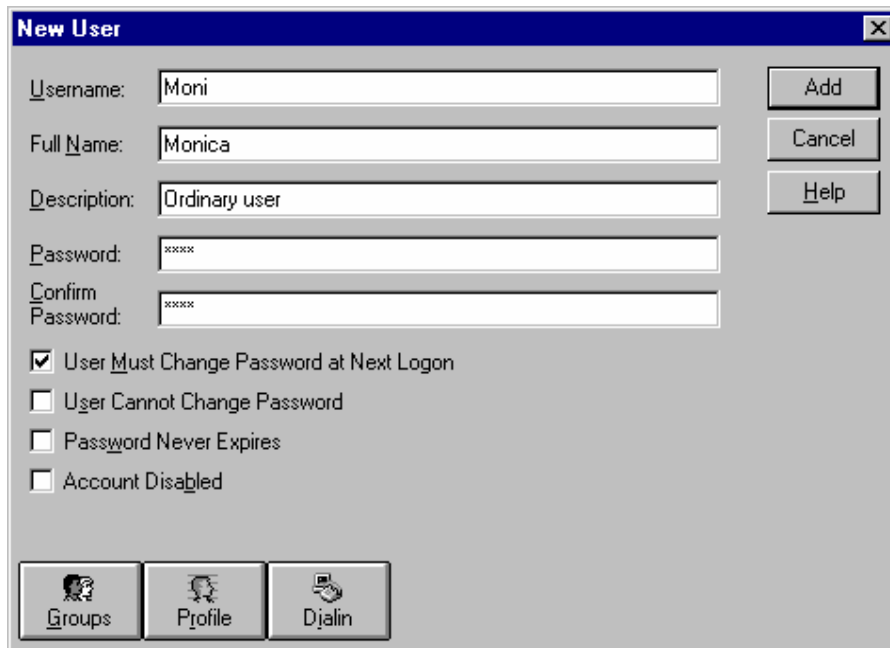


Figure 18.2 Creating a new account with the New User dialog.

The New User dialog contains many controls to which you must assign values:

- ❖ *Username.* Each network user must have a unique user name. The user name can contain up to 20 characters. You can use any combination of upper-and lowercase letters, numbers, or punctuation, except for the following characters:

= + [] / \ ; : < > ? * " ' `

- ❖ *Full Name.* Use this optional text box to enter the full descriptive name of the user for which the account is being created – for example, Fred Flintstone. As with the user name, the full name is better if you use a consistent method of full names, such as first name, last name, or last name, first name.
- ❖ *Description.* You can use the optional Description text box to further identify the user of this account – for example, a user’s department or title.
- ❖ *Password and Confirm Password.* The Password text box is used with the Username when users log on from a PC running Windows NT Workstation or Windows 95 with Windows NT Server authentication. You can leave the Password text box empty until users enter their password, but this leaves the network temporarily in a very unsecure state. If you leave the Password text box empty, be sure to select the User Must Change Password at Next Logon check box.

The Password text box is limited to 14 characters and is case-sensitive. It displays encrypted text as a row of asterisks (*). To make sure that the password is entered correctly, you must fill in the Confirm Password text box with a password identical to what’s in the Password text box before you can add the account.

- ❖ *User Must Change Password at Next Logon.* If this check box is marked, users must change their password the first time the account is used or the next time they log on to the domain. This check box, marked by default, should be used if users aren’t present when their accounts are being created or you can’t give them a password directly.

When you mark the User Must Change Password at Next Logon check box, the initial password must be either blank or something very intuitive to the user, such as the company name or the word *password*. This option allows users to customize their own passwords without administrative assistance. Be sure to verify that users immediately log on and change their password to avoid a potential security breach.

- ❖ *User Cannot Change Password.* The User Cannot Change Password option is primarily used when you are administering the passwords for user accounts centrally. This option is used primarily if several users share the same

account or in very secure networks. This option also is specified for the Internet Guest account, if present.

Select the User Cannot Change Password option only if users aren't allowed to enter their own password and you assign passwords to users.

- ❖ *Password Never Expires.* When this option is selected, users aren't required to change passwords periodically. Enabling this option isn't a good security practice, however, users should change their passwords on a regular basis, such as quarterly or even monthly. In certain cases – such as a rarely used account like the Administrator – you might use this option to avoid forgetting the password.

The Password Never Expires option is also used when security isn't a high priority compared to user convenience. Most users like to keep one password to avoid forgetting a new password. When this option is used, the check box User Must Change Password at Next Logon is cleared.

- ❖ *Account Disabled.* In certain instances, a user account must be disabled. Selecting this check box prevents a user from logging on to the network until the check box is cleared. Some of the reasons for disabling a user account are as follows:
 - Creating a template account that's used only to create new accounts by copying.
 - Disabling an account temporarily while a person is on vacation or extended leave.

Understanding the Additional Account Properties. When adding a new user, a set of six buttons appears at the bottom of the New User dialog: Groups, Profile, Hours, Logon To, Account, and Dialin. These buttons let you specify additional properties for user accounts.

Copying a User Account. To ease the task of setting up new user accounts, User Manager lets you copy an existing user account as a template to create a new account. In large networks, you create template accounts that contain all the attributes of a user in a particular department. When a new user account created, the appropriate template account is copied and the appropriate account information is changed to reflect the details pertinent to the new user.

To copy a user account, follow these steps:

1. From User Manager's window, select the user account to be copied.
2. Choose Copy from the User menu to display the Copy of Username dialog.
3. Enter the appropriate account information for the new user account.
4. Click the Add button to create the new user account.

Modifying User Accounts

User Manager lets you modify user accounts individually or modify multiple accounts simultaneously.

An individual user account can be modified by

- Double-clicking a user account item in the Username list of User Manager's window
- Selecting a user account in User Manager's window, and then either pressing Enter or choosing Properties from the User menu.

Either method displays the User Properties sheet, which looks similar to the New User dialog. The only significant difference between the User Properties sheet and the New User dialog is the addition of an Account Locked Out check box, which is used to clear a locked-out account. This check box is disabled unless the account is now locked out because of an excessive number of incorrect logon attempts.

You can modify multiple user accounts simultaneously by any one of the following methods:

- . Select a range of accounts by selecting a starting point in the list of user accounts and dragging the mouse pointer over the accounts to select.
- . Select a range of accounts by clicking a starting point in the list of user accounts and Shift + clicking the last account in the range.
- . Select multiple individual accounts by pressing the Ctrl key and then individually clicking each account to modify.

Choose **S**elect Users from the **U**ser menu to open the Select Users dialog. In this dialog, you can select or deselect all users assigned to a particular group within the domain you're administering. Multiple groups can't be simultaneously selected or deselected in this dialog.

After selecting the user accounts to be modified, press Enter or choose **P**roperties from the **U**ser menu to display the User Properties sheet .

When modifying multiple accounts, only the options common to all the users are displayed. The additional account property buttons located at the bottom of the dialog lets you assign common attributes to all selected user accounts.

Managing User Account Properties

Additional user account properties are accessed and managed through the buttons at the bottom of the New User and Copy of *Username* dialogs and the User Properties sheet: Groups, Profile, Hours, Logon To, Account, and Dialin. These property buttons let you specify additional properties for a user account. The following sections describe the dialogs that appear when you click each button.

Assigning Group Membership to a User Account

You assign group membership to a user account by clicking the Groups button. In the Group Memberships dialog, you can assign and revoke group membership privileges(figure 18.3).

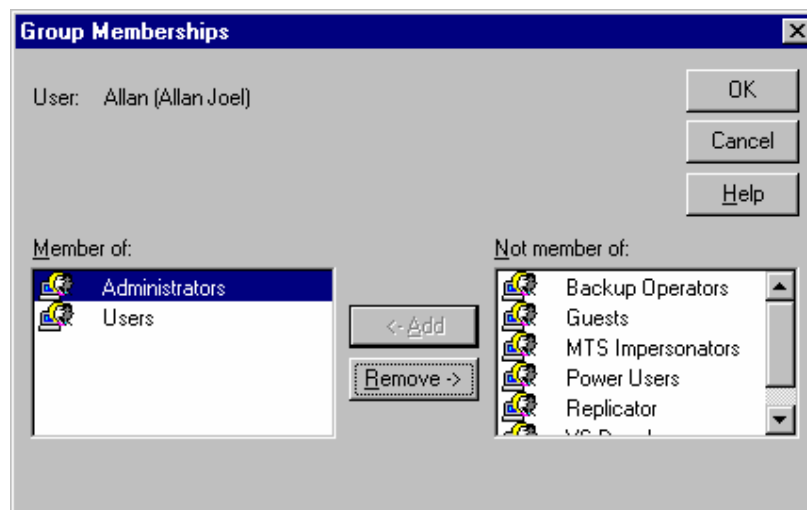


Figure 18.3 Assigning users to groups with the Group memberships dialog

The Member of list shows all the groups that the account belongs to. All groups to which the account does not belong appear in the Not Member Of list.

To assign a user to one or more groups, you can double-click the group to which to add the user account in the Not Member Of list, or select multiple groups in the Not Member Of list and click Add.

To remove a group membership from the user account, double-click the group to be removed from the account in the Member Of list, or select multiple groups in the Member Of list and click Remove.

In the Group Memberships dialog, the Set button applies only to users accessing a Windows NT network through Services for Macintosh. To set the primary group, select from the Member Of list and click Set.

Defining and Managing user profiles

To further define a user profile a user profile, use the Profile button to select custom settings for one or more users. Clicking this button displays the User Environment profile dialog.

User profiles provide power and flexibility for administrators and users when configuring a network environment. User profiles are typically stored in a common folder on a Windows NT server, Windows NT workstation clients also can have individual user profiles to supplement the user profile stored on the Windows NT server.

User profiles specify the startup information when individual users log on to windows NT. This information includes the user environment (environment variables, paths, and mapped drives), program groups, and available applications. When a user profile is stored on a central Windows NT server, the user environment is the same regardless of the computer the user logs on from. When the profile is stored on individual windows NT machines, the environment reflects the settings stored on each machine.

User profiles also can contain mandatory setting that you assign; users aren't allowed to alter these settings. This way, each user has a standard working environment and erroneous changes can't be made to a user profile, such as deleting a Start menu folder or shortcut.

Specifying the User Profile path. The user profile is specified in the User Profile path text box, which contains the location of a user profile located on a Windows NT server. If the text box is empty, profile is stored locally on each machine that the user logs on to.

Two types of user profiles are available;

- Personal user profiles. Each user is assigned this profile type and can alter the profile. Each user has his own private profile file with the file extension user.
- Mandatory user profiles. You assign this profile type, and the user can't change it. A mandatory user profile has the file extension man.

To assign a profile to a user account, type the profile path and file name in the User Profile Path text box. Be sure to follow the UNC naming convention. For example to store a profile named Profile.user on the computer RADI-SERVER in the folder Users, type \\radi-server\users\profile.user.

You can create the profile ahead of time by using the User Profile Manager. If the specified profile file doesn't exist when the user first logs, it's created automatically by using the default profile that exists on the workstation the user logs on from. Any changes are saved automatically to the profile file.

Setting a Logon script Name

Logon scripts are optional batch files that are run whenever a user logs on to a Windows NT network. Logon scripts are tailored to the client operating system that's used to log on to the network. All client operating system for Intel PCs except OS/2, use these .bat extension for the script file (OS/2 uses the .CMD extension). Windows NT Server 4.0 permits the user of executable logon scripts with an exe extension, but executable logon scripts are very uncommon.

Logon scripts aren't as flexible as user profiles but can be used instead of user profiles or with user profiles. By default, all logon scripts are stored in the folder \\Server Name\Winnt\system32\Rep1\Import\Scripts, where Server Name is the UNC server name of the primary domain controller for the domain you're administering. Because all scripts are stored in a central spot, only the name of the scripts file needs to be entered in the Logon Script Name text box. To use Server Manager to change the location of the Scripts folder, follow these steps.

1. Launch Server Manager from the Start menu's Administrative Tools (Common) choice.
2. Select the server in the opening window, and click the Replication button to open the Replication Properties for Server Name dialog.
3. Modify the entry in the Logon Script Path text box to point to your Scripts folder.

If a relative path, such as \Users\Logon.bat, is entered in the Logon Script Name text box for the user account, it's appended to the stored folder path. By using the preceding example, the logon scripts is run from the folder \\ServerName\Winnt\system32\Rep\Import\Scripts\Users\Logon.bat.

Logon scripts can be assigned on an individual basis, or the same logon script can be assigned to multiple users. Logon scripts for users of 16 bit windows must use DOS 8.3 file names.

Specifying a Home Folder. By default, a user is placed in the home folder when starting a DOS command session. The home folder also is used as a repository for user files and can be located on the client's local fixed or more commonly on a network drive.

To setup a home folder on a local machine, follow these steps.

1. In the User Environment Profile dialog (accessed by clicking the Profile button,) select the Local path option.
2. Type the local path (for example, c:\user\default) in the Local Path text box.

In a network environment, where a user can log on from multiple machines, the home folder should be located on a networked drive so that the user can access it from any machine. You can set up shared network drives for users to log on to. To set up a home folder on a networked drive, follow these steps:

1. In the User Environment Profile dialog select the Connect option.
2. From the drop down list box, select the drive letter of the client machine to contain home folder.
3. In the To text box, type a complete path for the home folder by using the UNC naming convention (for example, \\ radi-server\home\allan). This makes the home folder available to the user on any machine from which he logs on.

If the home doesn't exist, Windows NT creates it. Also the folder is protected so that only the specified user (and administrator) have access to the folder contents.

Managing Profiles for Multiple Users. When multiple users are selected in User Manger's window, the User Environment Profile dialog changes to reflect the selection of multiple user within groups (figure 18.4) . If all the user profiles to be modified are to share the same profile file name, logon script, and home folder, you make the same dialog entries as for individual accounts, as described in the preceding section.

You can streamline the process for creating individual user profiles based on a single profile for multiple users by using the environment variable %USERNAME%. Windows NT server 4.0 automatically replaces %USERNAME% with the user's logon ID. Assume that each user is to have an individual user profile with the file name derived from the user's first name. For the user names ALLAN,VIMALA and CHRISTY , supplying the path \\RADI-SERVER\Profiles\%USERNAME%.usr, has the same effect as creating three individual user profiles:\\RADI-SERVER\Profiles\ALLAN, \\RADI-SERVER\Profiles\VIMALA and \\RADI-SERVER\Profiles\CHRISTY. The variable %USERNAME% is expanded and replaced with the actual user name when any of the multiple users specified in the Users list of the User Environment Profile dialog log on to the Windows NT network. The %USERNAME% environmental variable can be used in any of the text boxed of the User Environment Profile dialog.

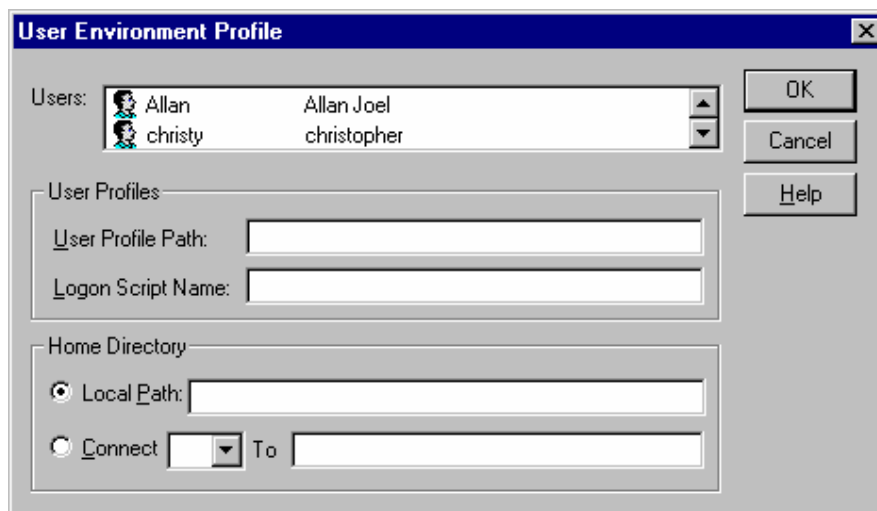


Figure 18.4 The User Environment Profile Dialog for multiple user accounts

Managing Logon Hours

When administering a large network, you might want to restrict the hours during which an account has access to the network. For example, certain workers may be able to access network resources only during normal business hours Monday through Friday from 8 a.m. to 5 p.m. Whereas other users have unrestricted network access. You manage logon hours for a user account by clicking the Hours button at the bottom of the New User and Copy of Username dialogs and the User Properties sheet to display the Logon Hours dialog .

The Logon Hours dialog displays a display schedule of times allowed for user logon. The dark areas indicate valid logon times. Logon hours are permitted by selecting the desired hours and clicking Allow. Similarly restricted hours are specified by selecting the hours and clicking Disallow.

You can use any of the following four methods to select logon times in the Logon Hours dialog.

- Clicking the day of week label for example, Sunday selects the entire day.
- Clicking the top of an hour column selects that hour every day of the week.
- Clicking the column square above Sunday selects the entire week.
- Clicking a specific hour selects that hour.

After the logon hours are set, click OK to save the logon hours for that account.

Setting Logon hours for multiple users. When managing logon hours for multiple users, select the desired users from User Manager windows and click the Hours button from the User Properties sheet. The Logon Hours dialog changes slightly from the single user version. If all the selected users don't have the same logon hours, a message box appears with the warning The selected users have different Logon Hours settings. If you continue the operation, all the user logon hours are reset and new logon hours are set in the same manner described earlier for setting the logon hours for an individual user account.

Logging off users who are logged on when logon hours expire. Although logon hours restrict when users can log on to a Windows NT network, users may be logged on when their logon time expires. The action that occurs in this situation is determined by the Account Policy set up by the domain administrator.

The following two actions can occur when a logged on user's logon time expires.

- Typically logged on users remain logged on, but they are denied the ability to make any new connections or access additional network resources.
- You can choose to forcibly disconnect logged on users on expiry of their specified logon times. When you choose this option, all logged on user receive a warning to log off the connected resource before the expiry time. Any users who don't log off before the logoff time are automatically disconnected.

Restricting logon privileges to assigned workstations.

To restrict which network clients users can log on to, click the Logon To button in the New User and Copy of Username dialogs or the User Properties sheet to open the Logon Workstations dialog.

By default, all new accounts unless the account is copied and the account that it was copied from has restricted access can log on from all clients. If client logon access needs to be restricted, follow these steps.

1. Select the User May Log On To These Workstation option.
2. Type up to eight client computer names in the text boxes.
3. Click OK to establish the restriction

When multiple users are selected in User Manager's windows, the Logon workstations, dialog displays information that applies only to all selected users. By selecting the option User May Log On To All workstations or by restricting client access, all selected users are affected.

Managing Account Information.

To assign specific user account information, click the Account button at the bottom of the New User and Copy of Username dialog or the User Properties sheet to open the Account information dialog. This dialog lets you determine the expiration date of the account and the type of account.

By default, a user account never expires. In situations where an expiration date is needed(such as when an employee leaves the company or for temporary employees), the account becomes inactive at the end of the day specified in the End Of date edit box.

The Account Type section specifies whether the account is global or local according to the following rules.

- Select Global account if the user account must be recognized by other domains that trust the user logon domain.
- Select Local Account if the user logs on to the domain from an untrusted domain or if user access is to be restricted to the logon domain.

When the Account information dialog is invoked for multiple users, only the properties common to all accounts are selected. Setting any of the options changes all selected accounts.

Setting user dial-in permissions

Windows NT Server 4.0 has eased the burden of granting dial-in permissions to users for Dial-Up Networking (DUN). In previous versions, dial-in permissions had to be assigned from the Remote Access Administration utility. Now you can assign dial-in permissions directly from User Manager by selecting the Dial-in button in the New User and Copy of Username dialogs or the User Properties sheet.

By default, users do not have dial-in permission; it must be granted to each user. Figure 18.5 shows the Dialin Information dialog.

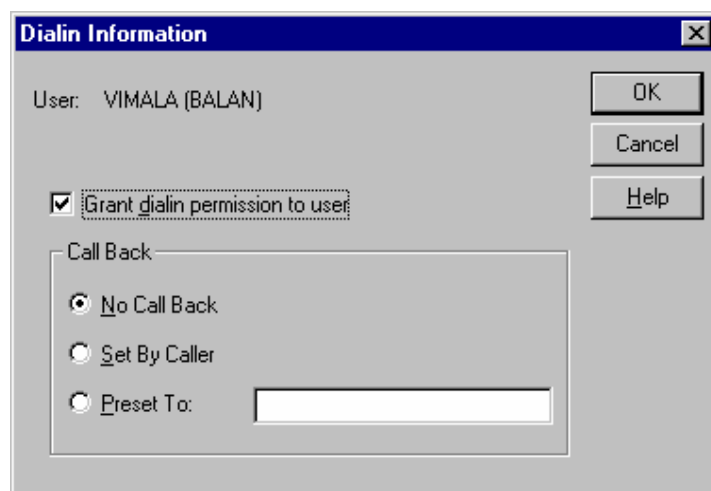


Figure 18.5 Setting callback properties in the Dialin Information dialog

If your network users need dial-in permission, simply mark the Grant Dialin Permission to User check box. Then set the Call Back options; No Call Back, Set by Caller, or Preset To.

If you select Set By Caller, users are prompted to enter an optional number that the server can use to call them back. This option is very valuable for users who travel a great deal, need to access the network for information while on the road, and want to minimize telephone charges.

The Preset To option is limited because the server called back to a specific number each time a user dials into the network. This option should be used only for strict security purposes where the user, usually a telecommuter, is always at a specific location.

18.4 Administering the Domain Account Policy

The domain account policy determines password and lockout restrictions for all users in the domain. Choose Account from User Manager's Policies menu to open the Account Policy dialog. The following sections describe how to set domain wide policies for passwords and account lockout.

Setting the Account Policy for Passwords

You can define the following types of password restrictions.

- *Maximum Password Age* determines how long account passwords are in effect before they expire. The options are password Never Expires or Expires In n Days (the default)
- *Minimum Password Age* determines how long a user account is forced to retain a new password. The objective is to prevent users from entering a dummy password when their password expires and then immediately changing it back to the old password. The options available are Allow Changes Immediately (the default) or Allow Changes In n Days.
- *Minimum Password Length* determines the minimum allowable length for all passwords. The options available are permit blank password or At Least n character. To maintain a secure network, require passwords of at least six (preferably eight) character.

- *Password Uniqueness* tells Windows NT server whether to keep a history of previously used passwords. The objective is to prevent users from reusing the same password when a password expires. The available options are Do Not Keep Password History (the default) or Remember n Passwords. For maximum security, set n to 8 or greater.

At the bottom of the Account Policy dialog is a Users Must Log On in Order to change password check box. If this options is selected and a user's password expires, the users of the account must ask the account administrator to change he password.

Setting the Account Lockout Policy

The Account Lockout Policy setting determines the actions that are taken if user forget their passwords or illegal attempts are made to access the network, as evidenced by multiple failed attempts to log on. In this event, either of the following actions can be chose;

- *No Account Lockout*. If this option is selected, any user can try an unlimited number of times to log on to the network.
- *Account Lockout*. If this option is selected, you set up lockout parameters to deter repeated illegal logon attempts. The following sections explain the lockout parameters.

Setting Account Lockout Options. One of the following two options applies to the Account Lockout setting.

- *Lockout After n Bad Logon Attempts* locks out the user account after so many failed logout attempts occur. This options forces the account user to wait until the account is unlocked through administrative or automatic intervention.
- *Rest Count After n Minutes* automatically resets the number of bad logon attempts to zero after so many minutes of account inactivity since the last bad logon attempt.

Setting lockout Duration options. One of the following two options applies to the Lockout Duration setting.

- *When forever (Until Admin Unlocks)* is selected, the account is locked out indefinitely until you manually reset the account.
- *When Duration n Minutes* is selected, the account automatically unlocks after so many minutes of locked time.

Removing users from the network when logon hours expire. When users are logged on to a windows NT network and their logon hours expire, you can either continue to let them access the network resources to which they are already logged on, or forcibly disconnect all users running windows NT workstation or Windows 95 from the network.

If the option Forcibly disconnect Remote User from Server When Logon Hours Expire is selected in the Account Policy dialog, remote users whose logon hours expire are prompted to disconnect from the network. If users don't log off the server will disconnect them automatically.

18.5 Short Summary

- A user account is created for each individual on the network and is meant for use only by that one person.
- Each user account includes the user name and password required to log on to the network.
- Account defines the user groups to which each user belongs and most importantly the right and permissions to access system resources granted to each user.
- User Account contains additional information, such as the user environment profile, a list of logon workstations and a schedule of logon hours.
- The Administrator account is the most powerful user on the network, having total access to and control over all resources within the domain for which the account is created
- The Guest account initially contains an empty password, which allows users from untrusted domains to log on to your domain as Guest and access any resources that are accessible to the Guest account.

18.6 Brain Storm

1. Define User account?
2. Explain the working of User Manager for Domains?
3. How you can manage the user account?
4. Write the various steps to create a New User Account?
5. How you can assign a group membership to a User Account?
6. Explain briefly about the Domain Account Policy?



Lecture 19

Administering Group Accounts

Objectives

In this Lecture you will learn the following:

- Understanding the concept of User Groups
- Able to manage the Group Accounts
- Understand the User Right Policy

Coverage Plan

Lecture 19

- 19.1 Snap Shot
- 19.2 Managing User Groups
- 19.3 Using Group Management Wizard
- 19.4 Managing User Rights Policy
- 19.5 Short Summary
- 19.6 Brain Storm

19.1 Snap Shot

This lecture described about the user groups , types of groups and how to add, copy and delete groups. The last section explain how to create new groups using Group Management Wizard. The following terms are basic to managing group accounts:

P5

- A *user group* is a management tool that collects user accounts into a named group. You assign rights and permissions to user groups in a manner similar to that for user accounts. You can grant user accounts membership in a user group, the account inherits all the rights and privileges of that group. The overall concept of Windows NT's user groups is similar to the group security implemented by database management systems such as Microsoft SQL Server and Microsoft Access.
- A *global group* is a collection of user accounts within a single domain. A global group can't contain other groups and can include only user accounts from the domain in which the group was created. This allow global groups to be assigned privileges anywhere on the network.
- A *local group* can be assigned privileges only in the domain in which the local group was created. Unlike a global group, a local group can contain users and global groups. Local groups let you collect groups from several domains and manage them as a single group in a local domain. When privileges are assigned to a local group, all users and global groups in the local group inherit these privileges.

19.2 Managing User Groups

The preceding sections of this chapter make many references to user groups. User groups define the rights and privileges that are assigned to the users in the groups. At the bottom of User Manager's window is a scrollable, alphabetically sorted list of the standard (built-in) groups of Windows NT server 4.0 (figure 19.1)

Two types of groups shown in the Groups list; global and local. A global group is depicted with a world globe in the background a local group is depicted with a workstation in the background.

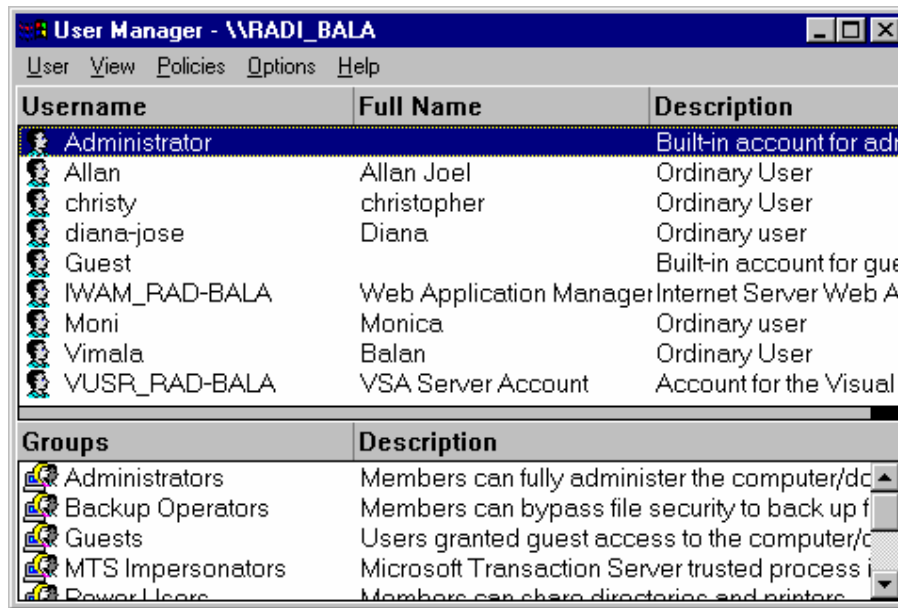


Figure 19.1 Built-in user groups available in User Manager Window

User Manager lets you create, modify, and delete groups assign user accounts to groups and remove user accounts from groups. The following sections describe the 11 built-in user groups of Windows NT server 4.0 and explain the management of user groups.

Examining Windows NT Server 4.0's Built-In Groups

The actions that a user account can perform depends on the group memberships assigned to the account, the rights and privileges the account inherits from the group(s) plus specific permissions you assign to the account. Windows NT Server 4.0 has 11 built-in user groups, each with a pre-established set of permissions for use of network resources. Descriptions of a few of these groups, by necessity, appear earlier in this chapter but are repeated here for completeness. Following is a brief description of each built in user group in the approximate order of decreasing privilege.

- The *Administrator* group is the most powerful local group in the domain. Administrators are responsible for the overall configuration of the domain and the domain's servers.
- The global *Domain Admins* group is a member of the Administrators group. By default, members of the Domain Admins group are as powerful as the Administrators group. The Domain Admins group can be removed from the Administrators group, in necessary to restrict the groups authority.
- The local *Users* group provides the capabilities that most users need to perform normal tasks. Members of this group have no rights to administer servers running Windows NT 4.0

- The global *Domain users* group is a member of the local Users group. By default, all new accounts are automatically added to this group, unless you specifically remove this group from the account.
- The local *Account Operators* group allows its members to utilize User Manager to create new groups and accounts. Members of this group have limited capabilities to administer account, servers, and groups in the domain. Group members can't modify or delete accounts or groups belonging to the Administrators, Domain Admins, Account Operators, Backup Operators, Print Operators, or Server Operators Groups, nor can they administer account policies.
- The local *Backup Operators* group can back up and restore files on the domain's primary and backup controllers. Members of this group also can log on to a server and shut it down, presumably for backup operations.
- The local *print operators* group allows members to create and manage printer shares in the domain. These members also can log on to a server and shut it down.
- The local *Server Operators* group allows members to manage the domain primary and backup controllers. Group members also can manage folder and print shares, as well as administer server functions, such as setting system time for the entire domain.
- The local *Replicator* group supports the capability to perform folder replication function. Only accounts needed to log on to the Replicator services of the primary and backup domain controllers should be members of this group.
- The global *Domain Guests* group is a member of the local Guests group. This group is intended for user accounts with more limited rights than a member of the Domain Users group.
- The local *Guests group* has very limited capabilities and is used for occasional or one time users.

Adding Local Groups.

The built-in user groups are adequate for most Windows NT Server 4.0 Networks. If you have a large, complex network, you might want to define your own user groups by, for example, organizational function or department. As an example, members of the Finance, marketing, sales and production departments might have their own groups. Similarly vice presidents, directors, managers and supervisors might be assigned to their own groups.

To add a local group to the domain, follow these steps:

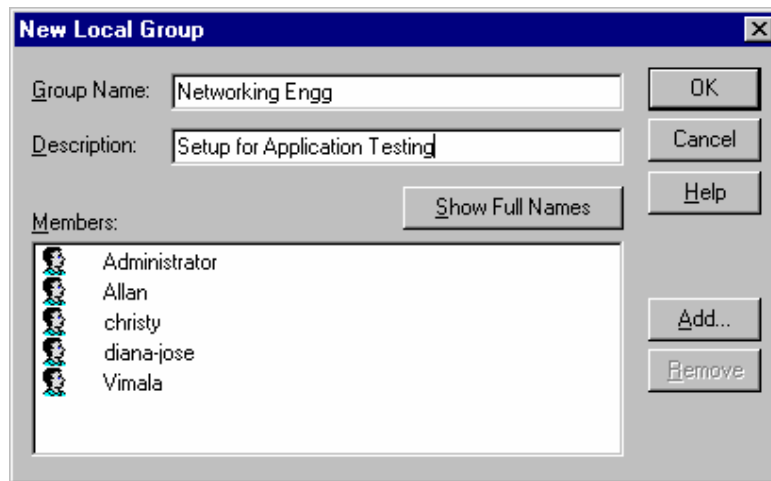


Figure 19.2 Adding a new local group with User Manager

From User Manager's User menu, choose New Local Group to display the New Local Group dialog (figure 19.2) .

1. In the Groups Name text box, type a group name that's longer than 20 characters. A Group name is required.
2. Type a group description in the Description text box. Although this is optional, a meaningful description is useful as your network grows and more groups are added.

To add user accounts to the new group, click the Add button to display the Add Users and Groups dialog. To add users to the local account, follow these steps.

1. Select the account entry in the Names list and click the Add button, or double-click the entry in the Names list to add the account to the Add Names list.(a description of each option in the Add Users and Groups dialog follows.)
2. Repeat step 1 for each additional account you want to add to the new group.
3. Click OK to add the accounts to the new group and close the dialog.

Local groups can include users and global groups from the domain of the local groups. Local groups also can include global users and global from other domains that are trusted by the local groups domain.

The purpose of the options in the Add Users and Groups dialog is as follows:

- *List Names From.* This drop-down list lets you select the domain from which to add names or groups. The default setting is the domain for the local groups.
- *Names.* This list displays all the users and global groups of the domain being viewed. The items in this list are candidates for inclusion in the new local group.
- *Members.* To view the members of a global group, select the global group from the Names list and click the Members button.
- *Search.* This button is used to find a domain name – a useful feature if your network contains many domains.

Adding Global Groups

The process for adding a new global group is identical to adding a local group, except that rather than choose New Local Group, you choose New Global Group from the User menu to display the New Global Group dialog. Unlike local groups, which can contain global groups and users, a global group can contain only users.

Copying a Group

If a new group needs to be created and will have similar rights and members as another group, it's easier to copy a group than to add a new group and manually set up its attributes. To copy a group, follow these steps:

1. Select a group to copy in User Manager's window.
2. From the User menu, choose Copy.
3. The Add New Local Group or Add New Global dialog appears, depending the type of group you selected in step 1.
4. Type a new name and description for the group.
5. Modify the group's membership, as necessary.
6. Click OK to create the new group and close the dialog.

Deleting Groups from the Domain

Only user-defined groups may be deleted from the domain. The built-in groups of Windows NT Server 4.0 can't be deleted.

Each group you create receives a unique security identifier (SID). If you delete a group and recreate a group with the identical name, the new group receives a different SID and doesn't inherit the original group's attributes.

To delete a group from a domain, follow these steps:

1. Select the group to delete from the Groups list of User Manager's window.
2. From the User menu, choose Delete. A warning message appears.
3. Click OK to proceed with the transaction, or click Cancel to abort the operation.
4. If you click OK, a second message asks the operator to confirm the decision. Click Yes to delete the group.

Deciding When to Use Local Groups or Global Groups

Determining when to add a local group or a global group to a domain can often be difficult. Use the following guidelines to determine whether to create a new global or local group:

- Use global groups when user accounts from this domain need to access resources of this domain and others.
- Use local groups when user accounts from this domain or others need to be used in resources of this domain. A local group should also be used when global groups from this domain or others need to be used in resources from this domain.

Providing Users In Trusted Domains Access to Resources Trusting Domains

Although one domain trusts another domain, the trust relationship doesn't grant users access to resources in the trusting domain. The easiest method for allowing users from other domains access to your resources is to add a global group from the outside domain to a local group in your domain.

User Manager lets you create global groups in other domains if the other domain trusts your domain. You can set up a global group in the external domain, select the user accounts needed from the outside domain and then assign that global group to a local group in your domain.

19.2 Using Group Management Wizard

The Group Management Wizard is a tool for creating new groups and adding users to the new group. You also can use the Group Management Wizard to change the membership of or delete existing groups.

The Group Management Wizard makes entering information into the New Global User or New Local User dialog a multistep process. Thus, it's questionable whether this wizard is of significant benefit to Windows NT network administrators. To decide for yourself whether use of the Group Management Wizard is worthwhile, follow these steps:

1. From the Start menu, choose Programs, Administrative Tools, and Administrative Wizards to open the wizard selection dialog; then double-click the Group Management icon to display the first Group Management Wizard dialog.
2. The wizard lets you create a new group and add members or modify the membership roster of an existing group. To create a new group, accept the default Create a New Group and Add Members option. Click the Next button.
3. Type the name of the new group and an optional description in the two text boxes. Click Next.
4. If you're working at the server on which the group is to be created, accept the default option. On My Computer; other wise, select On Another Computer. Click Next.
5. If you select a computer that's a domain controller, a message box tells you that the new group will be created on the primary domain controller. Click OK to continue.
6. You can choose between creating a Global Group (the default) and a Local Group. Unless you have a specific reason for creating a Local Group, accept the default and click Next.
7. All users appear in the Available Members list. Select each user you want to join to the new group and click Add to add the user to Selected Members list. When you've added all the members, click Next.
8. The last dialog confirms the name and domain for the new group. Click Finish to add the new group to the domain.
9. A message confirms addition of the new group. If you've had enough group management wizardry for the moment, click No. If you want to give the Group Management Wizard another try, click Yes.

19.3 Managing User Right Policy

Each user's capabilities are determined by the rights and privileges assigned to the user. A user's *rights* refer to the entire system or domain. User Manager assigns all rights. The rights assigned to a user directly affect the tasks that user can perform on the network.

Determining User Rights

A Windows NT network has two categories of rights: basic and advanced. Below Table lists the user rights of Windows NT Server 4.0 and the built-in groups that receive these rights.

Basic Rights for the Built-in Windows NT Groups

User Right	Group Rights Assigned To...
Access this computer from network	Administrators, EVERYONE
Add workstations to domain	Administrators, Backup Operators, Server Operators.
Backup files and directories	Administrators, Backup Operators, Server Operators.
Change the system time	Administrators, Server Operators
Force shutdown from a remote system	Administrators, Backup Operators, Server Operators
Load and unload device drivers	Administrators.
Log on locally	Account Operators, Administrators, Backup Operators, Print Operators, Server Operators.
Manage auditing and security log	Administrators.
Restore files and directories	Administrators, Backup Operators, Server Operators.
Shut down the system	Account Operators, Administrators, Backup Operators, Server Operators.
Take ownership of files or other objects	Administrators
Bypass traverse checking	EVERYONE
Log on as a service	Replicators
Assign user rights	Administrators
Create and manage local groups	Administrators, Users
Create and manage user accounts	Administrators
Create common groups	Administrators
Format computer's hard disk	Administrators
Keep local profile	Administrators, EVERYONE
Lock the computer	Administrators, EVERYONE
Manage auditing of system events	Administrators
Override the lock of the computer	Administrators
Share and stop sharing directories	Administrators
Share and stop sharing printers	Administrators

Assigning New User Rights

When you create a new user group, you can customize the rights assigned to the group by adding or deleting them. To add or delete rights from group membership, follow these steps:

1. In User Manager's window, choose User from the Policy menu to display the User Rights Policy dialog.
2. The Right drop-down list at the top of the dialog displays rights that you can assign to or remove from Windows NT Server 4.0 groups. By default, only basic rights are listed. To see advanced rights, mark the Show Advanced User Rights check box at the bottom of the dialog, and then select the right you want to examine.
3. To add new groups to the right, click the Add button to display the Add Users and Groups dialog with a list of users and groups in the domain. Select the groups and users that you want to assign the right to .
4. When all your changes are complete, click OK to effect the changes and close the dialog.

19.5 Short Summary

- A user group is a management tool that collects user accounts into a named group.
- A global group is a collection of user accounts within a single domain.
- A local group can be assigned privileges only in the domain in which the local group was created.
- Local groups collect groups from several domains and manage them as a single group in a local domain
- Permission assigned to a user refer to the specific files, folder and hardware devices accessible to that user.
- The term EVERYONE isn't a group, but a Windows NT convention for indicating that all users in all groups have the given rights
- If you choose to modify an existing group, you select the computer on which the group was created and then select the group to modify. You can delete the group or in a succeeding dialog, modify the membership of the group

19.5 Brain Storm

1. Define Local and Global groups ?
2. How you can examine Windows NT Server 4.0's Built-In Groups ?
3. What are the steps to add a local and global groups to a domain?
4. How you can copy and delete the groups from the domain?
5. How you can create a new groups using group Management Wizard?



Lecture 20

Sharing & Securing Network Resources

Objectives

In this Lecture you will learn the following:

- Able to share and secure Folders and Files
- Able to compress NTFS Files and Folders
- Knowing the concept of Replication folder

Coverage Plan

Lecture 20

20.1 Snap Shot

20.2 Sharing & Securing Folders and Files

20.3 Replication Folders

20.4 Sharing and Securing Network Printers

20.5 Short Summary

20.6 Brain Storm

20.1 Snap Shot

In this lecture you are going to learn how you can share and secure the Network Resources. The fundamental purpose of any network operating system(NOS) is to give users access to shared network resources such as folder, files and printers. Just as important as the capability to share these resources is the capability to control which users have access to each resource.

Windows NT provides all the tools you need to share and secure folders and files. You can control access to folders and files on a very broad level. For example, folder shares of servers that use the FAT file system function like a blunt-instrument. Folder shares let you share a FAT folder on the Windows NT server but allow access control only at the group level and only then to the folder and all subfolders as a group.

Windows NT server also provides all the tools you need to share and secure network printers. You can share printers that are physically connected to the computer running Windows NT Server. You also can share printers physically connected to other Microsoft Networking clients on the Network, configuring them to appear as shared resources on the Windows NT server.

20.2 Sharing and Securing Folders and Files

Windows NT Server makes sharing folders and files easy for you. Behind this ease of use lurks the power needed to control which users can access which resources. In the following sections, you learn how to share folders and files and how to control access to them. The extent to which you can secure your folders and files depends on the file system you decide to use.

Choosing a Windows NT Server File System

During its history, Windows NT has supported the following three file systems:

- ❖ The FAT *file system* is marginally faster than the other file systems on small servers but provides none of the data integrity features available with the HPFS and NTFS file systems. Access control is limited to share-level security. Don't consider using the FAT file system on a production server.
- ❖ The *High Performance File System(HPFS)*, originally by IBM for its OS/2 operating system, is fast and provides good data integrity features, but offers only share level security. However, NTFS includes all the features of HPFS, so there's no reason to use HPFS. Unlike prior versions of Windows NT 4.0 doesn't support

HPFS, although you can access HPFS files and folders running on networked Windows NT 3.x servers.

- ❖ The *NT File System (NTFS)* is the native file system designed by Microsoft for Windows NT Server. It's fast, offers excellent security, and provides rock-solid data integrity functions.

Although Microsoft offers you a choice of file systems, don't spend too long thinking about which to pick. Using NTFS provides the best mix of speed, security, and protection for your data. You can compress files and folders stored only on NTFS partitions. Even if you take Microsoft's recommendation to use FAT for your system partition, use NTFS partition(s) for all server shares.

Understanding Folder Shares

Until a folder is shared, only administrators can access it across the network. Administrators have access to default administrative shares for each logical server drive and all its folders. Administrative shares, are create automatically by Windows NT Server.

Folder shares provide the first level of security by controlling which folders on the server are visible to – and therefore accessible by logged-on users. As a means of securing access, folder shares have the following drawbacks:

- Sharing a folder automatically shares all files contained in that folder and its subfolders. If you need finer control of which subfolders and files are accessible to which users, you must use folder-access and file-access permissions, which are available only if you're using the NTFS file system.
- A folder share controls access only for those users who log on to the server from a remote workstation. Any user with physical access to the server can log on locally and bypass share-level security.

Creating, Modifying, and Removing Folder Shares. To create a folder share, you must be logged on locally to the computer running Windows NT Server, and your account must be a member of the Administrators, Server Operators, or Power Users group. Follow these steps to create a new folder share:

1. Open Explorer or double-click the My Computer icon to display a list of drives available on your server.
2. Double-click one of the available drives to display a list of folders contained on that drive. If the folder you want to share isn't at the root level, click the + symbol to the left of the parent folder name to display a list of subfolders for that folder.
3. Right-click a folder to display the context-sensitive menu.
4. Choose Sharing to display the Sharing page of the property sheet for that folder.
5. By default, the folder is marked Not Shared. Select the Shared As option button to activate the remaining controls of the dialog so that you can enter information for the share.
6. Type a descriptive name for the share into the Share Name combo box. This is the name by which users access the shared folder. Optionally, type a more complete description of the resource into the Comment text box.
7. Specify User Limit information. By default, the new share is set to Maximum Allowed, which allows any number of users to access the share simultaneously, up to the limit of the number of users for which the server is licensed.

Select the Allow option button and select a specific number of allowable simultaneous users, if you want to limit the number of users permitted to access this share at any one time. Do this if you're concerned about performance degradation when a large number of users view for a single resource.

By default, the new share provides Full Control to the group Everyone. This means that any user with an account on the server can add, modify, or delete files contained in this folder. The following section, "Working with Share Permissions," describes how to restrict access to the new share.

To remove a folder share, perform the preceding first four steps to display the Sharing page of the *Foldername* Properties sheet select the Not Shared option button and then click the Apply button.

To modify the share, specify a new Share Name, Comment, or User Limit, as described in the preceding steps. You also can create an alias for this shared resource by clicking the New Share button and completing the dialog. An *alias* allows the same shared resource to be accessed by more than one share name.

By clicking the Permissions button, you can determine which users and groups have access to this shared resource, and at what level. The following section describes how to restrict access in this manner.

Working with Share Permissions. Share permissions control which users and groups can access a share, and at what level. You can add, modify, view, or remove the following share permissions for each folder you have shared on the server.

- *No Access* permission restricts all access to the shared folder.
- *Read* permission allows users to view file names and subfolder names within the shared folder. You can change to a subfolder and open a file in the shared folder or in a subfolder in read-only mode, but you can't write to that file or delete it. You can execute program files for which you have only Read permission.
- *Change* permission grants all the rights provided by Read permission and adds the rights to create new files and subfolders, modify the contents of new or existing files, and delete files and subfolders.
- *Full Control* permission grants all rights provided by Change permission and adds the rights to create and modify NTFS file permissions and folder permissions, as well as take ownership of NTFS files and folders.

Follow these steps to modify, view, and remove share permissions:

1. Perform the first four steps in the preceding section to display the Sharing page of the *Foldername* Properties sheet.
2. Click the Permissions button to display the Access Through Share Permissions dialog. The Name list displays the users and groups authorized to access this share. By default, the group Everyone is assigned the Full Control permission to the share.

3. To modify the share permission for an existing user or group, select that user or group and then select a Type of Access from the drop-down list.

To remove the share permission for an existing user or group, select that user or group and click the Remove button.

4. Click OK to accept the changes and return to the *Foldername* Properties sheet.

Adding a share permission requires a few more steps. Follow steps 1 and 2 in the preceding list to display the Add Through Share Permissions dialog. Then proceed as follows:

1. Click the Add button to display the Add Users and Groups dialog.
2. Select the domain or computer from which the new users or groups are to be added from the List Names from drop-down list. Member groups of the selected domain or computer are displayed in the Name list.
3. Select a displayed group by clicking its name. (By default, only groups are displayed. To display users, click the Show Users button.)
4. After you select all users and groups to be added to the share, use the Type of Access drop-down list to select the access type to be granted to the selected users and groups.
5. Click the Add button and then click OK to add the selected users and groups to the share. The Access Through Share Permissions dialog reappears, with the new users and groups added to the share and their access type displayed.
6. In the Access Through Share Permissions dialog, click OK to return to the *Foldername* Properties sheet. Click OK to accept the changes you've made to the share.

Using Administrative Shares. In addition to the shares you create, Windows NT Server automatically creates several shares for administrative purposes. These administrative shares include at least the following:

- ADMIN\$ points to the location of the shared Windows NT Server folder on the server. For example, if you install Windows NT Server to the C:\Winnt folder on your server, the ADMIN\$ share points to this folder.
- [drive letter] \$ points to the root folder of each drive on the server. For example, if your server has three drives, designated C, D and E, these drives are each represented by administrative share, named C\$, D\$, and E\$, respectively.

The most common administrative shares are the drive and folder shares. However, administrative shares can also represent a named pipe for remote procedure calls (RPCs), a communication-device queue (only on LAN Manager servers), or a shared printer.

Displaying All Shares and Disconnecting Shares. A distinctive icon in Windows NT Explorer and the My Computer window indicates shared folders. However, sometimes seeing a comprehensive list of shares displayed in one place is useful. To see a list of all active shares on your server, proceed as follows:

1. From Control panel, double-click the Server icon to display the Server dialog.
2. Click the Shares button to display the Shared Resources dialog. For each share, this dialog displays the Sharename, Users (the number of current active sessions for the share), and Path associated with the share name.
3. To disconnect one share, select its name and click the Disconnect button. To disconnect all shares in one step, click the Disconnect All button.

Using the Managing Folder and File Access Wizard

The Managing Folder and File Access Wizard provides a quick and easy way to create and manage folder shares. To use the Managing Folder and File Access Wizard, follow these steps:

1. From the Start menu, choose Programs, Administrative Tools, and Administrative Wizards to display the Administrative Wizards window. Click the Managing File and Folder Access icon to display the first dialog of the Managing Folder and File Access Wizard.

2. Select On My Computer to create or manage shares on the server, or select On Another Computer to manage shares on another computer on the network. In this example, a new share is created on another server. Click Next to display the dialog.
3. Select the computer where you want to create or manage the share and click Next. The Managing Folder and File Access Wizard displays the dialog. In the example, a new share name is entered into the To Create a New Folder. Type a New Name text box to create a new share named SHARED.
4. Click Next to display the Managing Folder and File Access confirmation message. Click Yes to create the new folder. The Managing Folder and File Access Wizard displays another message box to confirm that the new folder has been created successfully. Click OK.
5. Click Next to display the next Managing Folder and File Access Wizard dialog, which allows you to set permissions for the folder to determine who has access to it, and at what level. By default, the original permissions for the share are retained, and these permissions flow down to affect the files and subfolders contained within this folder.

To change these default permissions, click Change Permissions and choose one of the three options presented:

- Only I Have Access and Full Control
- I Have Access and Full Control, Everyone Else Can Only Read It.
- Everyone Has Access and Full Control

Mark the Apply These Permissions to All Folders and Files Within this Folder check box if you want the permissions you set here to apply to all subfolders and files contained within this folder. Unmark the check box if you want these permissions to apply only to this folder.

6. Click Next to display a Managing Folder and File Access message box, which lets you specify whether the folder will be shared with network users. Click Yes to allow network users to access the folder.

7. The Managing Folder and File Access Wizard displays the dialog. You can rename the share, provide a brief description of it, and specify which types of network users can access it. Make any changes necessary and click Next.
8. The Managing Folder and File Access Wizard displays the summary. Click Finish to complete creating the share.
9. The final message box lets you exit the Managing Folder and File Access Wizard or continue managing shares. Click No to exit or Yes to manage another share.

20.3 Replication Folders

Windows NT Server4.0 allows you to replicate or copy, folders to other computers or domains to maintain identical copies of folders and files on more than one computer. The folder from which data is copied is called the export folder and is located on the export server, the folder to which data is copied is called the import folder and is located on the import computer. The export and import folders can be located on the same computer or on different computers.

Folder replication does more than simply copy data from the export folder source to the import folder destination. The Windows NT server replication service functions much like an FTP mirror program. It monitors the export folder for changes to existing files and newly created files and subfolders and replicates these changes and additions to the import folder. The replication service also deletes files in the import folder that have been deleted from the export folder, By doing so , it synchronizes the contents of the two folders

Folder replication is most commonly used for the following two purposes:

- ❑ *Replication logon scripts from one domain controller to other domain controllers* allows users of any domain controller to log on locally and reduces server load and network traffic.
- ❑ *Replicating a database from one server to another* allows users who access the database to be distributed among two or more server in order to share the work load among multiple servers.

You also can use folder replication to keep a frequently updated backup copy of a heavily used database file, which would otherwise be difficult to back up.

Creating a Replication User.

Before you can configure the replication service you must first create a special user for that service as described . This new special user must have the following properties:

- The user must be assigned to the backup Operators group
- The Password Never Expires check box must be marked.
- The Logon Hours settings must allow this user access at all times.

You won't be able to name the new user Replicator because a group already exists with that name. Choose another similar name, such as Replicate.

Starting the Replication Service.

After you create the special user, you must then configure and start the Directory Replicator service before folder replication can occur. To do so, Proceed as follows:

- 1 From Control Panel , double-click the Services tool to display the Services dialog. The status is shown as blank, indicating that the Directory Replicator service isn't running. Start up is shown as Manual, indicating that this service won't be started unless you do so manually.
- 2 With the Directory Replicator service selected , click the startup to display the Service dialog.
- 3 In the Start up Type Section, select Automatic to indicate that the Directory Replicator service start automatically each time Windows NT server is started.
- 4 In the Log on As section, Select this Account, and then enter the domain and user account name that you created in the preceding section. You can also click the ... button to display a list of available accounts to choose from .

Type the password for this account in the password and confirm password text boxes.

- 5 Click OK to accept the changes. You're prompted to restart Windows NT Server.
- 6 After Windows NT Server is restarted, double-click Control Panel's services tool to verify that the Directory Replicator service has been started successfully.

Configuring Folder Replication

After you successfully configure the directory Replicator service, you must then configure an export server and an import computer.

To configure the export server, you must provide the following pieces of information.

- ❑ The export folder designates the source folder from which files and subfolders are exported.
- ❑ The export to list designates computers and domains to which files and subfolders are exported. If you designate a domain here, exported data is replicated on all computers in the export- to domain that have replication enabled.

To configure the import computer, you must also provide two pieces of information, as follows

- ❑ The import folder designates the destination folder in which imported files and subfolders are stored.
- ❑ The import from list designates computers and domains from which data to be imported is accepted.

To configure the export server and the import computer, proceed as follows:

- 1 From control panel, double-click the Server tool to display the Server dialog.
- 2 Click the Replication button to display the Directory Replication dialog.

- 3 In the export section. Select Export Directories to enable exporting. Then designate which folder is to be exported in the From Path text box. Click the add button to add domain or computers to the To List to designate a target or targets to which files are exported.
- 4 Click the Manage button to display the Manage Exported Directories dialog. You can use the controls in this dialog to add and remove exported directories and to add and remove locks on managed directories.
- 5 If this server will also be an import computer, select the Import Directories option in the import section of the Directory Replicator dialog to enable importing. Then designate which folder to receive the imported data in the To Path text box. Click the Add button to add domains or computers to the Form List to designate computers and domains from which imported data is to be accepted.
- 6 Click the Manage button in the import section to display the Manage Imported Directories dialog. You can use the controls in this dialog to add and remove imported directories and to add and remove locks on managed directories.

20.4 Sharing and Securing Network Printers

Beyond sharing folders and files the most common purpose of most networks is to share printers. One justification for early local area network was their capability to share expensive laser printers among many users. In the Past few years, the prices of laser printers have plummeted; It's now economically feasible for many companies to provide sub-\$1,000 personal laser printers such as the Hewlett-Packard Laser Jet 5L and 5P to any client that needs one.

Still in all, the original justification for sharing expensive printers on the network holds true. Ten years ago, your might have been sharing a \$3,500 LaserJet that printer eight letter-size pages per minute at 300 dpi. Today you might instead be sharing a laser printer that prints 2011-by-17 inch pages per minute at 600 dpi, but that printer still costs \$3,500 and budget realities still demand that it be shared. Just as it always did, the network allows you to share and expensive resources such as high speed laser printers and color printers.

Windows NT Server makes it easy to share printers on the network. Printers attached directly to the Windows NT Server can be shared as a network resource and used by any network client authorized to do so. Network clients running Windows 3.11 for Workgroups, Windows 95, or Windows NT Workstation can also function as printers servers, sharing their attached printers with other network users.

Configuring Locally Attached Server Printers as Shared Resources

After you physically install the printer to be shared and connect it to the Windows NT Server, you can use the Add Printer Wizard to configure it and make it available as a shared printer. To do so, proceed as follows:

- 1 From the My Computer window, double-click the Printers icon to display available printers (If you haven't installed any printers yet, only the Add Printer icon appears in the Printers Window).
- 2 Double-click the Add Printer icon to invoke the Add Printer Wizard. You can select My Computer to add a printer to the local computer or Network Printer Server to add a network printer that's physically connected to a different computer section describes adding a locally connected printer, so select My Computer and then Next.
- 3 The next dialog lets you specify the port to which the printers connected, add a port, and modify the properties for a port. Mark the check box that corresponds to the port that your new printer is connected to.
- 4 If you need to add a port to the available Ports list, click the Add Port button to display list of available printer ports. When you add a printer port and accept the change by clicking OK you return to the preceding Add Printer Wizard dialog where the newly added printer port appears as an available selection.
- 5 In the Second Add Printer Wizard dialog, you can click the configure Port button to display and modify Port settings. If the selected port is a parallel port, the configuration Port dialog opens.

If the selected port is a serial port (also called a COM port), the Ports dialog opens. Select the COM port that the printer is connected to and click the

Settings button to display the Settings for COMx dialog. Select the settings for Baud Rate, Data Bits, Parity, Stop Bits and Flow Control from the drop-down lists that correspond to the settings of the printer being installed.

- 6 Click the Advanced button to display the Advanced Settings for COMx dialog. In this dialog, you can adjust settings for COM Port Number, Base I/O Port Address and Interrupt Request Line (IRQ). The FIFO Enabled check box , when marked allows Windows NT to use the buffering provided by 16550 and higher UARTs (Universal Asynchronous Receiver/Transmitter) to improve Windows printing performance. If an advanced UART was detected during Windows installation this check box is marked by default and should be left marked. If Windows NT didn't detect and advanced UART on this port during installation the check box is disabled.
- 7 After you finish selecting the printer port, click OK to advance to the Add Printer Wizard printer selection dialog. Select the manufacturer of your printer in the Manufacturers list; then, in the printers list, select the model of your printer. Click next.
- 8 The fourth Add Printer Wizard dialog lets you specify whether this printer is shared, to provide a share name for the printer and to load support for other operating systems that will be printing to this printer. Complete this dialog and Click Next. If you've specified that support for operating system other than Windows NT 4.0 is to be loaded, you're prompted to insert driver disks for those operating systems.
- 9 The next Add Printer Wizard dialog lets you print a test page. You should always allow the wizard to print the page to verify that your printer has been installed successfully and is performing as expected. After you print the page and verify that it printed correctly, click the Finish button.
- 10 The Copying --Files needed dialog prompts you to insert the Windows NT Server CD-ROM so that the necessary files can be copied from it. Specify the drive and path name for these files or click the Browse button to browse the location. Make sure that the CD-ROM disk is inserted in the drive, and click OK proceed with copying files.
- 11 When all needed files are copied from the Windows NT CD-ROM , the Add Printer Wizard prompts you to insert the distribution media for the other

operating system you've elected to provide printing support for. Insert the media and specify the location of these files as described in step 10.

The Add Printer Wizard now takes you directly to the Printer Properties sheet so you can configure the newly installed printer. This process, used both to configure newly installed printers and to reconfigure already installed printers, is described in the following section.

Configuring Network Printer Servers as Shared Resources

The preceding section described how to configure a printer that's physically attached to the computer running Windows NT Server as a shared Printer. The Add Printer Wizard also configure a network printer server as a shared resource on the server, A network printer server is a print queue that services a printer that's physically connected to a different computer on the network.

In this section, you learn how to configure a printer queue serviced by a Novell NetWare printer server as a Windows NT Server shared resource. You can use the same procedure to associate a Windows networking printer queue with a share name on your Windows NT server, letting the present printers connected to Windows Networking clients as a server shared resource.

To install and configure a network printer server as a shared server resource, follow these steps:

- 1 From the My Computer Window Double-click the Printers icon to display available printers.
- 2 Double-click the Add Printer icon in the Printers Window to invoke the Add Printer Wizard. You can select my Computer to add Printer to the local computer(as described in the preceding section), or Network Printer Server to add a network printer that's physically connected to a different computer. The section describes adding a network printer server, so select Network Printer Server and click the Next button.
- 3 The connect to Printer dialog opens, displaying the available networks and network printer queues visible to Windows NT Server. Double-click the

- printer server name to display the print queue is associated with that printer server. Click OK.
- 4 If the selected print queue doesn't have a printer driver installed, you're prompted to install an appropriate driver locally on the Windows NT server. Click OK to install the driver locally.
 - 5 The Add Printer Wizard moves next to selecting a printer manufacturer and model. Select the manufacturer of your printer in the Manufacturers list and then select the model of your printer in the Printers list. Click the Next button.
 - 6 The Connect to Printer -- Copying Files -- Files needed dialog prompts you to insert the Window NT Server CD-ROM so that the necessary files can be copied. Specify the drive and path name for these files, or click the Browse button to browse for the locations. Make sure that the CD-ROM disk is inserted in the drive and click OK to proceed with copying files.
 - 7 When the necessary files are copied the Printer properties sheet opens . The exact contents of this dialog vary, depending on the capabilities of the particular printer you're installing. Configure these settings appropriately and then click OK.
 - 8 The Add printer Wizard default printer dialog asks you whether this printer should be set as the default printer. Select the appropriate option and click Next.
 - 9 The final Add Printer Wizard dialog opens. Click Finish to complete installation of your network print queue printer and return to the Printers window.

Configuring Printer Properties

The following procedure is automatically invoked as the final step in installing a local printer, described. When use in this fashion, the Add Printer Wizard places you at step 3 in the following procedures. You also can use this procedure to reconfigure an existing printer, beginning with step 1:

- 1 From the My Computer window, double click the printers icon to display available printers.
- 2 Select the printer you want to configure in the printers window and right-click to display the context-sensitive menu. Choose Properties to display the General page of the Printername Properties sheet.
- 3 On the General page, supply the following information:
 - ❑ *Comment* lets you enter a short comment that can be viewed by users of the printer. For example, if the printer is available only during normal business hours note that in this text box.
 - ❑ *Location* lets user view the physical location of the printer to make sure that they know where to pick up their print jobs.
 - ❑ *Driver* lets you select from a drop-down list of available drivers for the printer
 - ❑ *Click the new Driver* button to install a new or updated driver for the printer.
 - ❑ *Click the Separator Page* button to specify options for separator paves, used to keep print jobs separate.
 - ❑ *Click the Print Processor* button to select different methods of processing the incoming byte stream. Using the default Win Print processor unless you have specified reasons for changing it.
 - ❑ *Click the Print Test page* button to print a test page to verify printer functioning.
- 4 After you complete the General page, display the Ports page. You can use the Add Port, Delete Port, Configure Port buttons to modify the port configuration of the printer as described in the preceding section. The Enable Bidirectional Support check box is marked by default if your printer supports this function; if it doesn't this selection is grayed out.
- 5 After you finish configuring the port, display the scheduling page. This page lets you specify when the printer is available to users, at what priority

print jobs are to be handled, and the various options to control how spooled documents are processed.

The following options are available from the scheduling page:

- ❑ *Available* defaults to Always, allowing users to access this printer at any hour. You can select the From option and specify From and To times if you want to restrict availability of the printer to specified hours.
- ❑ *Priority* lets you specify what priority level windows NT Server assigns to this printer.
- ❑ *Spool Print Documents so Program Finishes Printing Faster* lets you specify that incoming print jobs are written to a temporary file and processed from that file. If you select this option, you can choose between start Printing After last Page is Spooled and Start Printing Immediately. In the first case, Windows NT Server waits until the entire print job has been written to a temporary spool file before it begins printing the document. In the latter case, windows NT Server begins printing as soon as it has received enough data to complete the first page. The later selection is marked by default because Start printing Immediately almost always provide better printing performance. If your network is very heavily loaded , you may need to specify Start Printing After The Last Page is Spooled to prevent page from different print jobs from being interleaved, and other printing problems.
- ❑ *Print Directly to the Printer* Lets you specify that incoming print jobs are sent directly to the printer without first being queued.
- ❑ *Hold Mismatched Documents*, marked retains documents in the queue that couldn't printed successfully because of mismatched pages
- ❑ *Print Spooled Documents First*, if marked gives preference to printing documents contained in the Spool before other documents.
- ❑ *Keep Documents After They Have Printed*, if marked, retains documents in the print queue even after they print successfully. Windows NT Server ordinarily removes documents from the print spool after they're printed. Marking this check box results in all documents being retained in the

spool, which causes a rapid growth in disk space consumed for spooled documents. Mark this check box only as a part of diagnosing printing problems.

- 6 After you finish setting scheduling options, display the Sharing page . The upper section of this page lets you specify that the printer be Not shared or shared. If it's a set as shared you can modify the share name in the share name text box.

The bottom section of the Sharing page lets you specify alternate drivers that allow users of other operating systems to use the shared printer. The alternate Drivers list shows that support is installed only for Windows NT 4.0 running on the X86 processor family. You can install support for additional operating systems by selecting them in this list. Later, when you finally accept changes to all pages of the Printername Properties sheet by clicking OK, you're prompted to insert the disks containing the printer drivers needed.

- 7 After you finish setting sharing options, display the Security page which has three sections each of which is accessed by clicking that sections button. The permission button lets you specify which groups are permitted to access the printer. The Auditing button lets you specify by user and by group which actions are recorded to an audit log. The Ownership button lets you specify which user or group owns the printer.
- 8 Click the permissions button display the Printer Permissions dialog. The name list displays the name of each group that's now authorized to access the printer on the left , with that group's level of access specified on the right. You can add a selected group by clicking the Add button and responding to the prompts. You can remove a selected group by clicking the Remove button.

To change the access level associated with a selected group or groups select the type of access to be allowed from the Type of Access drop-down list. You can assign one of the following types of access:

- No Access* allows the group so assigned no access what ever to the printer

- ❑ *Print* allows the group so assigned to print documents, but not to manage the printer or modify its properties. You should assign this access level to ordinary users of the printer.
- ❑ *Manage Documents* allows the group so assigned to print documents and manage the printers. Manage documents is normally assigned to the creator /owner of the printer.
- ❑ *Full Control* allows the group so assigned to print documents manage the printer and modify its properties. Full Control should normally be assigned to the group administrator, Print Operators and Server Operators.

After you set permissions as necessary, click OK to return to the security page of the *Printername* Properties sheet.

- 9 In the Security page of the *Printername* Properties sheet, click the Auditing button to display the Printer Auditing dialog. By default, no auditing is assigned for the printer. To add auditing for specified users and groups, click the add button to display the Add Users and Groups dialog. To add users and groups to the Add Names list, either double click the user or group name or select the name and click the Add button.

After you finish adding users and groups, click OK to return to the printer Auditing dialog. Domain users added for auditing as well as auditing configured to report only Print Failure for the selected group. After you specify the desired level of auditing for each selected group, click OK to accept the changes and return to the security page to the *Printername* Properties sheet.

- 10 In the Security page of the *Printername* properties sheet, click the Ownership button to display the Owner dialog. You can take ownership of this printer by clicking the Take Ownership button, or close the dialog by clicking close. In either case, you return to the Security page of the *Printername* Properties sheet.

- 11 In this *Printername* Properties sheet, display the Device Settings page. The appearance of this page varies, depending on the characteristics of the printer for which you're setting properties. After you configure the device settings to your satisfaction, click OK to save the property settings for all pages.

20.5 Short Summary

- To add several users and groups to share in a single step, select multiple users and groups by using standard Windows selection conventions.
- If you're using share permissions to restrict access to a shared folder, remember to remove the default share permission that grants the group Everyone the Full Control share permission for that folder.
- Share permissions are cumulative, so any user has all share permissions granted to any group of which he is a member
- Share permissions specify the maximum level of access available within the shared folder tree.
- Any networking server or client can share an attached printer as a network resource.

20.6 Brain Storm

- 1 Explain how a folder and files can be shared?
- 2 Write the steps to create , modify and remove a share folder?
- 3 Write the steps to modify, view and remove share permissions?
- 4 How you can display and disconnect all shares?
- 5 What is the purpose of using replication folder and how you can create a replication user?
- 6 How you can share and secure the network printers ?
- 7 Write the steps to install and configure a network printer server as a shared server resource?



Lecture 21

Monitoring the Network

Objectives

In this Lecture you will learn the following:

- Understand the concept of Performance Monitor
- Able to Monitoring and examining disk performance
- Knowing the network statistic logs
- Understand the concept of Event Viewer and types of Event

Coverage Plan

Lecture 21

- 21.1 Snap Shot
- 21.2 Performance Monitor
- 21.3 Network Monitor
- 21.4 Event viewer and Log files
- 21.5 Short Summary
- 21.6 Brain Storm

21.1 Snap Shot

An important part of network management involves monitoring trends on the network. By effectively monitoring network behavior, you can anticipate problems and correct them before they disrupt the network. Monitoring the network also provides you with a baseline, a sampling of how the network functions in its equilibrium state. By establishing a baseline on your system, you can determine whether your network can handle the current resource usage or whether additional resources are needed.

This lecture presents various programs or mechanisms that can be used to monitor and record information about the network. The explanation of what these different mechanisms are and when you would utilize them is addressed in the lecture.

Monitoring Network Trends

Monitoring the network is an ongoing task that requires data from several different areas. The purpose of this lecture is to bring these tools together so that you can view them in the context of an overall network monitoring strategy. The following list details some tools you can use to document network activities:

- Written documentation
- A statistics- gathering or performance-monitoring tool such as Windows NT's Performance Monitor
- A network- monitoring and protocol analysis program- such as Windows NT's Network Monitor or the more powerful Network Monitor tool included with Microsoft's BackOffice System Management Server(SMS) package – or a hardware based protocol analyzer.
- A system event log, such as the Windows NT event log, which you can access through Windows NT's Event Viewer Application.

21.2 Performance Monitor

The Windows NT Performance Monitor is a tool that enables you to pick apart individual components on a computer (including hardware and software), record statistics, graph those statistics and view real-time statistics. Right out of the box Performance Monitor offers more than 350 statistics for your Windows NT system.

Figure shows the Performance Monitor application. You start this application by choosing it for the administrative tools program folder from the Start menu or by running the command PERFORMS.

When you start Performance Monitor, the application opens by default into what is known as Graph View. This is the view that is used the most and that enables you to view statistics in real-time.

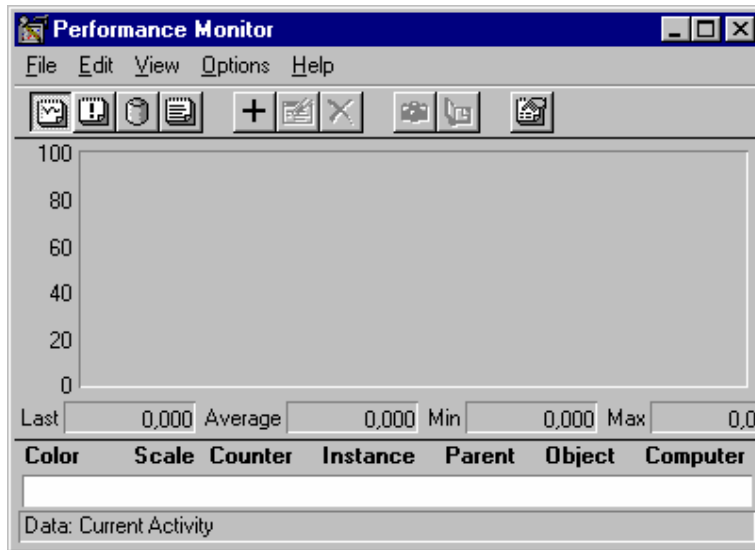


Figure 21.1 *The Windows NT Performance Monitor*

Objects and Counters in Performance Monitor

Performance Monitor can monitor many statistics. Performance monitor refers to the entity that is being monitored as an object and the statistic that it captures as a counter. Windows NT has internal counters that monitor activity on the computer, and Task Manger can display these counters. Performance Monitor uses the same counters, but can access a greater depth of these counters. The following table lists some of the most important objects of Windows NT, the counters associated with each object and the significance of each counter.

Counter	Description	Significance
Memory		
Available Bytes	Amount of virtual memory available	When the value falls below a threshold, Win NT gradually takes memory from running applications to maintain a certain minimum of

		available virtual memory.
Pages/Sec	Number of pages that had to be written to or read from disk and placed in physical memory	This counter indicates whether more physical memory is needed in your system.
Page Faults/Sec	Number of page faults in the Processor	This value indicates that the data needed wasn't immediately available on the specified working set in the memory.
Paging File		
% Usage	Amount of the paging file (Pagefile.sys) in use	This indicates whether you should increase the size of Pagefile.sys.
% Usage Peak	Peak usage of system paging	This indicates whether the paging file is of appropriate size.
Avg. Disk Sec/Transfer	Average no of bytes transferred to or from disk during read/ write operations	Low values of this counter indicate that applications are accessing the disk inefficiently.
Avg. Disk Sec/Transfer	Amount of time of disk takes to fulfill requests	A high value can indicate that the disk controller is continually Retrying the disk because of read or writes failures. A high value is greater than 0.3 second.
Disk Queue Length	No of disk requests outstanding at the time performance data is collected	This value related to the number of spindles that make up the physical disk. A single disk has one spindle. RAID drives have multiple spindles but appear as single drive. A typical value is up to two times the number of spindles making up the physical disk.
% Disk Time	Percent of time that disk is in use.	If the value consistently is above 85% active, consider moving some files to an additional server or upgrading the disk drive.
Processor		
% Processor Time	Percentage of elapsed time that processor is	This counter indicates how busy a processor is. If this value is very

	busy executing a non-idle thread	high, the system may benefit from a processor upgrade or multiple processors.
Interrupts/Sec	Rate of service requests from I/O devices	This indicates the number of requests processed from device drivers. If this value increases without corresponding increase in system activity, it could indicate a hardware problem on the system.

This is only a small subset of the objects that can be monitored by Performance Monitor. And the number of objects that Performance Monitor can monitor is extensible. All the Microsoft BackOffice products will add objects to Performance Monitor. Microsoft has provided an application programming interface (API) to enable developers to add objects to Performance Monitor that relate to their applications.

One of PerMon's most important capabilities is its ability to monitor the performance of Windows NT objects on remote machines. For example from a client running Windows NT Workstation 4.0, a system administrator can monitor the performance of objects on all Windows NT server in a domain. This feature is very useful for detecting load balancing problem in the network.

The monitor behavior of a remote computer's objects by following these steps.

1. If PerMon isn't displaying a chart, choose Chart from the View menu.
2. Choose Add to Chart from the Edit menu to open the Add to chart dialog.
3. In the Add to Chart dialog's computer text box type the name of the Windows NT computer to monitor. Alternatively, click the button to the right of the text box you display the Select Computer dialog, which display a list of all computer on the network.

If you use the Select Computer dialog, double click the name of the domain or workgroup to display a list of servers and workstations in the entity. Select the name of the server or workstation to monitor, and then click OK to close the dialog.

4. In the Add to Chart dialog, a new chart is added to PerMon by selecting an object from the Object drop down list, selecting a counter for the object from the Counter list, and then clicking the Add button. The processor object and % Processor Time counter selected.

5. Click Cancel to close the Add to Chart dialog. The line chart begins to display % Processor time

Charting Performance Characteristic

Performance Monitor typically creates line charts, but histograms or bar charts are alternative for certain types of data. When using charts, PerMon displays the collected statistical data near the bottom of the windows in the value bar. The value bar shows the following information.

- *Last* shows the most recent reading that was taken.
- *Average* display the average of all data readings taken since the chart was created.
- *Min* shows the lowest reading that was taken.
- *Max* displays the highest reading that was taken.
- *Graph* time shows the amount of time that one chart on screen covers.

Creating a new chart. Creating a new chart requires that you select a computer, object, and counter to monitor. The following example shows how to create a new chart with the Processor object and % Processor Time counter. Monitoring the percentage of processor time consumed indicates whether the server's CPU is overworked. Such a condition indicates that the server needs a faster processor or that some of the work performed by this server should be distributed to another server.

To create a new line chart for monitoring % Processor Time, follow these steps.

1. Choose Chart from the View menu or click the View a Chart button button on the toolbar. (if a chart already exists, you can clear it by choosing New Chart from the File menu.)
2. Add to the chart a new line representing the object being monitored by choosing Add to chart from the Edit menu or by clicking the add Counter button in the toolbar. The Add to Chart dialog appears. (you must complete the rest of the steps before the new line appears in PerMon)
3. In the Computer text box, type the name of the system to monitor, or click the button to the right of the text box to open the Select Computer dialog, which shows all computers on the network.
4. In the Add to Chart dialog, select the object you want to monitor from the Object drop down list. (for example, choose the Processor object) . When you select an object, the Counter list fills with the appropriate counter for that object.

5. Select the counter that you want to monitor from the Counter list. In this example select the % Processor Time counter.
6. If the computer has more than one CPU, the Instance list box fills with an instance for each CPU. Select the instance to monitor. In this example, select 0 for the first CPU.
7. You can choose the Color of the object, the Scale factor, the line width, and the line Style. These options enable you to identify more distinctly each counter when the same chart is monitoring more than one counter.
8. Click the Add button to add the new performance characteristic to be monitored to the chart.
9. You can add more performance characteristics to the chart by repeating steps 3 through 8.
10. After you add all the objects that you want to monitor, click the Cancel button to return to the Performance Monitor window. PerMon now begins monitoring the chosen objects.

Editing a chart. When your chart is running, you can edit any counters being monitored. As shown earlier in figure 21.1 legend of all monitored counter appears as the bottom of PerMon's windows. To edit any counter, select the counter to be edited from PerMon's legend at the bottom of the window. Then from the Edit menu choose Edit Chart Line, or double click the counter in PerMon's legend to open the Edit Chart Line dialog. The Edit Chart Line dialog resembles the Add to Chart but doesn't allow for Computer, Object, Counter, or Instance changes. The only changes you can make are to the Color, Scale, Width, and Style of the line representing the object edited. Make the appropriate changes and click OK.

Deleting Monitored Objects a Chart Deleting objects from the chart often becomes necessary during the course of monitoring system performance. You may need to delete an object due to inconclusive results obtained from it, or simply to clear a chart that's monitoring many performance characteristics. To delete any counter beings monitored, select it on PerMon's legend and press the Delete key.

Customizing the chart options. You can modify the presentation of PerMon charts by using the options in the chart Options dialog. To open the dialog, choose chart, from the Options menu. The following table describes all the available chart options.

Option	Description
Legend	When this default option is selected the legends for each chart line appear at the bottom of PerfMon.
Value Bar	When this default option is selected, the value bar shows the values of Last, Average, Min, Max and Graph Time.
Gallery	The Gallery setting determines how the data is displayed. The options are Graph, which is the default and most useful method, or Histogram, which displays data as a bar graph.
Update Time	The default Periodic Update setting, the most commonly used method, tells PerfMon to read new data at the time interval specified in the Interval (Seconds) text box, which defaults to one second. Manual Update tells PerfMon to update on user request.
Vertical Grid	This options display grid lines on the vertical axis.
Horizontal Grid	This option display grid lines on the horizontal axis.
Vertical Labels	This default option displays labels for the vertical (y) axis.
Vertical Maximum	This option specifies the maximum value for the vertical axis. This value should be changed to reflect the counter being monitored. For example, set vertical maximum to 100 when using a percentage counter.

Using Performance Monitor to Set Alerts

One useful feature of PerfMon is the capability to define alerts, which can be sent to any station on the network when a monitored counter reaches a critical value. An alert is a method of monitoring any counter and performing a specified action when the counter exceeds or falls below a predetermined threshold value. You can choose to log the alert, send a notification message to a user on the network, or run an application.

You can view the alerts defined by performance monitor by either choosing alert from the view menu or clicking the View the Alerts button of Performance Monitor's toolbar.

Adding an Alert You add alerts to PerfMon by performing steps similar to the steps used for adding charts to PerfMon. Creating a new alert removes any existing charts that PerfMon is displaying. To add a new alert to the view, follow these steps.

1. Choose Alert from the View menu, or click the alert toolbar button.
2. Choose Add to alert from the Edit menu, or click the add button, to open the Add to Alert dialog.

3. Type the name of the system to monitor in the computer text box, or click the ... button to open the select computer dialog, which shows all computers on the network.
4. From the object drop-down list, select the object that you want to monitor, for this example, choose the processor object.
5. Select the counter that you want to monitor from the counter list box. In this case, select the % processor Time Counter.
6. If the computer has more than one CPU, the instance list box fills with an instance for each CPU. Select the instance to monitor. In this example, select instance 0 for the first CPU.
7. Select a color for the alert.
8. Enter the alert threshold in the Alert If section, and specify whether the alert is to run if the value is under the threshold or Over the threshold.
9. When the alert triggers and if an application is supposed to run, specify the application name in the Run Program on Alert-text-box. Indicate whether the alerts is to be run the First Time the alert triggers or Every Time the alert triggers.
10. When the options are set, click the Add button to add the alert to PerfMon.

Customizing alerts As with charts, you can customize alerts with various options. Choosing Alert from the Options menu opens the alert options dialog. The below table describe the available alert options.

<i>Alert option</i>	<i>Description</i>
<i>Switch to Alert View</i>	PerfMon switches to alert view when an alert is triggered.
<i>Log Event in application log</i>	PerfMon generates an entry in the application log when the alert is triggered. You can view the application log from Event monitor.
<i>Send Network Message</i>	PerfMon sends a network message to the machine specified in the Net Name text box. Don't enter double backslashes(\\) before the machine name.
<i>Update Time</i>	Two options are available here. The most used is periodic update, which tells PerfMon to read new data at a time interval specified in seconds in the interval text box. Manual update tells PerfMon to update on user request.

PerfMon uses NetBIOS as the transport protocol for network messages. If your alerts are sending network messages, you must ensure that NetBIOS is available as a transport protocol and that the NetBIOS messenger service must be running for alerts to send network messages to the intended recipients.

You can determine whether the NetBIOS messenger service is started by using control Panel's Services tool. The service name Messenger appears in the list box of all services, along with the current status of the service. If this service doesn't appear, you must install the NetBIOS interface protocol. From the service page of control Panel's Network tool.

The following two methods start the messenger service:

- From the command prompt, type `net start messenger`.
- In control panel, double-click the services tool. Then select the messenger service from the list box displaying all services and click the start button. You can change the startup properties of the service to have it start automatically when the computer boots.

After the messenger service starts, make sure that the alert recipient is added by typing `net name machinename add` at the command prompt, where *machinename* is the name you typed in the Net name text box in the Alert options dialog.

Using Performance monitor log files

You use log files to provide a history of how your network is operating. You can set up PerfMon to keep a log for the results of running charts and of alerts that occur. Maintaining log files can ease the burden of network administrators as the network grows and performance begins to degrade. You can examine log files to determine the source and location of bottlenecks and devise plans for correcting the problems.

Recording data to a log file To record data to a log file, follow these steps:

1. Select the log view by choosing Log from the View menu or by clicking the view output log file status toolbar button. PerfMon displays a new log view window.
2. From the File menu, choose New Log settings to clear existing log settings.

3. A new log file is added by choosing Add to log from the Edit menu or by clicking the add toolbar button. The add to log dialog appears.
4. Type the name of the system to monitor in the computer text box, or click the ... button to open the select computer dialog.
5. From the objects list box, select the object that you want to monitor. For this example, choose processor. Unlike in chart view, when an object is selected, all instances of the object are logged.
6. After you select the object you want to log, click the add button to add the object to the log.
7. To log multiple objects, repeat steps 5 and 6 for each object. Then click cancel to close the dialog.
8. Now that the objects to log are selected, the log file needs to be set through the log options dialog, which you can open by choosing Log from the Options menu.
9. Enter a new log file name, or select an existing log file to overwrite from the log options.
10. In the update Time section, either select periodic update and set a time interval in seconds, or select manual update for user intervention. These selections determine when data is written to the log file.
11. Click the start log button to start logging data. PerfMon's window changes to the log view.

To stop data logging, open the log options dialog and click the stop log button.

Viewing Recorded data logged data, unlike chart data, doesn't appear in PerfMon's window you can open for viewing only log files that aren't currently opened. If the log you want to view is in use, you must first stop the log by choosing Log from the view menu. Then choose log from the options menu to open the log options dialog and click stop log.

To view the log data, follow these steps.

1. From the options menu, choosing data from to open the data from dialog.
2. Select the log file option.
3. Type the path and name of the log file to view in the text box, or click the ... button to find the log file.

4. Click OK to close the data from dialog.
5. To specify a time frame to view within the log file, choose time window from the edit menu to open the input log file timeframe dialog. The bar above the bookmarks section indicates the time line for logging events in the file. The bookmarks section lists events when data was logged, the default is the entire time for the log file.
6. To set a new start time, select a bookmark and click the set as start button.
7. To set a new end time, select a bookmark and click the set as stop button.
8. Click OK to close the input log file. Timeframe dialog and display the logged data.

To resume viewing current activity, you must open the data from dialog and select the current activity option.

21.3 Network Monitor

The main purpose of network monitor are to allow you to see what traffic is passing on your network at the protocol and packet levels. As such, it's not a tool for the casual user, and well only cover the rudiments of using it here.

There are a large number of widgets, buttons, graphs, and panes, displaying a smorgasbord of data. Each network monitor capture window (so called because you use it to capture frames of network traffic) represents a single network interface card (NIC) if you have more than one NIC, you'll have more than one capture window within the network monitor window. Each capture window is titled with the adapter's device name (for example \ethernet\net 1) and contains four panes. In the upper left corner is the graph pane, which displays bar graphs of usage and timing parameters. Immediately below it is the session stats pane, which displays statistics for the currently active network session. The total stats pane takes up the right hand side next to the graph and session stats panes, it displays a summary of network and capture statistics. The bottom of the window is ordinarily taken up by the station status pane, which shows overall statistics for the computer begin monitored.

Any of the panes can be shown or hidden. All four panes are shown in the default Configuration, but you can hide or show any of the panes with the window menu or the second group of buttons on the toolbar.

You can customize the panes display to suit you each of the columns in the session stats and station stats panes can be resized and you can reorder the columns by dragging them with the left mouse button into the order you want. You can also customize how network monitor displays captured data. One final note in either the capture window or

the summary window you can save the current view settings fonts, colors, columns widths and so on with the capture or display menu's save configuration command.

Installing network monitor and the network monitor agent The NTS installation doesn't default to installing the network monitor or network monitor agent, nor does it provide you a way to ask for it during the install. If you want to monitor network traffic on you server, or if you want any of you NTW or NTS machines to be monitored by an SMS server running the SMS network monitor, you need to install the network monitor package on the machines to be monitored.

There are actually two components you can install , the network monitor agent component, which has only the network monitor agent software, or the network monitor tools and agent component, which includes the network monitor application with the network monitor agent software. If you want to monitor an NTS or NTW machine, install the network monitor tools and agent component instead.

In either case, the steps needed to add network-monitoring capability are similar to those required for installing SNMP?

1. Open the network control panel applet and choose the services.
2. If the menu shows "network monitor agent" or "network monitor tools and agent" the network monitor agent has already been installed, so skip on to step 4.
3. Click the add button when the select Network service dialog box appears, select which package you want to install on this machine, then click OK. NT setup will prompt you for the path to the installation media, then it will copy the files you've selected. After the installation is finished, you will be reminded to restart before the new components become active.
4. To start the network monitor agent service, open the services control panel, applet and look for the network monitor agent entry. If you just installed network monitor, the service will be listed as a manual startup service. Click the startup button, then set the startup type to automatic and click OK. Finally click the start button to start the network monitor agent.

Capturing Network Data

Before you can analyze traffic you must capture it. Network monitor's capture menu furnished commands for configuring how capturing works, from how much data can be captured to what data is filtered during the actual capture. Once captured, data can be written to disk and filtered or analyzed later.

Network monitor copies data from the network card's onboard buffer to its own RAM buffer. The size of this buffer and the speed of your computer determine how much data you can capture and how much of the data from the wire you actually capture. The capture buffer settings command brings up the capture buffer settings dialog box, with which you can specify how much RAM will be used for the capture buffer as well as how much of each incoming frame to capture. Use the buffer size field to allocate a chunk of RAM for the network copy buffer.

You can increase the effectiveness of your buffer by telling network monitor to capture only the part of the network frames that interest you. By default, network monitor captures the entire frame, but you can use the frame size field to capture only the first N bytes of the frame, skipping over the application specific contents if you don't need them.

Since network data can arrive at speeds up to 100Mbps, your computer may fail to capture some packets if it can't keep up. Network monitor offers two ways to capture data. Normal mode keeps the statistics panes in the window up to date as data arrives, and dedicated mode freezes the statistics panes and minimizes the network monitor application to maximize your capture rate. To toggle between the two options use the capture dedicated capture mode command.

Setting Capture Triggers

Network monitor allows you to set triggers that will automatically take some action when the condition you set becomes true. For example you can set a trigger to stop capturing data when a particular pattern occurs. You might do this to get details on an FTP sessions that keeps failing to certain hosts, setting a capture trigger that stops the capture when the session completes you to capture only the protocol data of interest.

You set capture triggers with the capture trigger dialog box, which appears when you choose the capture trigger command.

The trigger on group offers you five choices for when the trigger should occur.

- The default is nothing, which means nothing will happen. (the other choices are more interesting)
- Pattern Match causes a trigger event when the pattern you specify in the pattern group occurs in a frame.
- Buffer space causes a trigger event when the percentage of the buffer in use hits the value you specify. You're limited to choosing space in 25% increments.

- Pattern match then buffer space links a pattern with a buffer space limit. Network monitor will first watch for the specified pattern when it occurs, it will fire a trigger when the buffer utilization reaches the value you set.
- Buffer space then pattern match makes network monitor wait until the buffers filled to the point you to choose when it does, the trigger will fire when the specified pattern occurs in an incoming frame.

Once you've hit on the trigger condition that suits your needs, you can set a trigger action to be executed when the trigger is fired. Network monitor can either do nothing (the default) or stop the capture. In either case, you can specify that an external program be run by checking the execute checkbox and specifying the command line for the program in the command line field.

Starting, Pausing, and Stopping Capture

Once, you've constructed a capture filter that will keep only the data that interests you, you're ready to capture data. Network monitor provides a simple mechanism for logging network data. When you're ready to start, use the capture start command (or the start captures toolbar button). If you're in normal capture mode, the statistics panes update in real time to show you what's happening, if you're in dedicated mode, network monitor minimizes itself and displays a small dialog box with buttons for pausing or stopping the capture and a counter indicating the number of captured frames.

While you've capturing data, you can pause and restart the capture at any time with the capture pause and capture continue commands (or the pause/continue capture toolbar button,) When in dedicated mode, use the buttons in the summary dialog box.

Viewing and Analyzing Captured Data

Once you've stopped capturing data, you can move on to viewing and analyzing the packets that were captured by your filter. You can start viewing data in two ways, by selecting the capture stop and view command when you stop capturing or the capture display captured data command at any time after you've stopped the capture. In either case, network monitor will open a new summary window to show you what data was captured.

Network Monitor's user interface changes somewhat when you bring a summary window frontmost. The capture menu changes its name to display, and show some additional options. An edit menu appears; it can be used to copy data from the window and paste it into a text editor or other application.

Property and protocol options are added to the help menu(Which protocols have help available depends on the network vendor; the standard network monitor installation only includes help for Microsoft's SMB protocol).

The standard summary window will list a single entry for each captured frame. You can expand an entry to see more details of the frame by double-clicking an entry in the summary. The summary pane will shrink and two new panes will be added to the window. The middle pane, which shows you detailed data about the frame you selected., the bottom pane, the Hex pane, shows you the actual contents of the packets. Items in the Details pane may have sub items you can collapse or expand these items by double clicking them or single clicking the plus or minus sign in the left most column.

While you're looking at the frame summaries and details, you can move between frames with the display menu's next frame, previous frame, and go to frame commands. Goto frame requires you to pick a frame number, but the others go to the selected frame relative to the current frame. You can also use the find next frame command to find frames that match your selection criteria, these criteria are identical to the ones you'll construct for display filters.

The summary window display can be customized, too. Like the session stats and station stats panes of the capture windows, you can adjust column widths and ordering by dragging the column rules or titles with the left mouse button. You can rename the capture window, too. The default name for the summary window is the name of the capture buffer (the name you used to save the capture buffer or "capture" if it hasn't been saved.) You can rename the capture window's with the window label command. Whatever label you type it will appear in the window's title bar when it's visible.

Network monitor also lets you customize the font and color used to display the protocol information. Each protocol can have its own foreground and background colors, but the font you choose is used for all protocols. Use the display font and display colors commands to set these the way you like.

21.4 Events Viewer and Log Files

Event Viewer provides information about such events as errors, warnings and the success or failure of tasks. An event is any potentially significant occurrence in the system or in an application. Some critical events, such as a full drive are logged. Non critical events are merely logged. Event logging starts automatically each time Windows NT is started. With an Event log and Event Viewer, it is possible to troubleshoot various hardware and software problems and to monitor Windows NT Security event

To start the Event Viewer click the Start button, point to Programs, points to Administrative Tools and then click Event Viewer

Types of Events

Error Significant problems such as loss of data or of functions for example, an Error event is logged if a service is not loaded during Windows NT startup

Warning Events that are not necessarily significant, but that indicate possible future problems. For example, a Warning event is logged if disk space is low

Information Infrequent but significant events that describe successful operations of drivers or services. For example when a network driver loads successfully, it logs an information event

Success Audit Audit security access attempts that are successful. For example, a user's successful attempt to log on to the system is logged as a Success Audit event

Failure Audit Audited security access attempts the fail. For example, if a user tried to access a network drive and fails, the attempt is logged as a Failure Audit event

Event Log Files

Event log files are records of system events. Windows NT records events in three kinds of log the System log, the Security log, and the application log. The system and Application logs are automatically enabled: the Security log must be manually enabled

Types of Log

System (systemroot\System32\Config\Sysevent.evt) contains events logged by the Windows NT system components and device drivers. For example, the failure of a driver or other system component to load during startup is recorded here. The event types logged by system components are determined by Windows NT, and those logged by third-party drivers are determined by the driver vendor

Security (systemroot\System32\Config\Secevent.evt) can contain valid and invalid logon attempts, as well as events related to resource use, such as creating, opening, or deleting files or other objects.

Application (systemroot\System32\Config\Appevent.evt) Contains events logged by applications. For example, a database program might record a file error here. Application developers decide which events to monitor

Interpreting an Event

View events stored in the log files in greater detail by clicking Detail on the View menu in Event Viewer. The Event Detail dialog box shown the following information:

- Date and Time of the event
- Event identification
- Text description of the selected event



In some cases, the Event Detail dialog box displays hexadecimal data for the selected event. The component that was the source of the event record generates this information. If interpretation of the hexadecimal data is required to resolve a problem, a person who is familiar with the source component should be contacted. Not all events generate such hexadecimal data.

There are several methods for controlling the Event display and finding specific events in Event Viewer. These methods including Filtering, arranging and searching.

Filtering

When Event Viewer starts, all events recorded in the selected log are displayed automatically. To view events with specific characteristics on the View menu click Filter Events

Events can be filtered by using the properties shown below:

View From/ View Through From a specified date and time, or during a period of time

Types Error, Warning, Information, Success Audit and Failure Audit

Source The software that logged the event. An application or component of the system such as the Browser or a Floppy disk

Classification Defined by the source. For example, the categories for Security event logs include logon, logoff, and policy change and object access

User Specific text that exactly matches the text in the User name field

Computer The exact name for the computer on which the logged event occurred

Event ID A number used to identify the specific event

Arranging

Events displayed in Event Viewer are listed in sequence, from the, most recent to the oldest. By using the View menu the sequence can be changed to display the oldest events first.

Searching

To search for events in large logs, use the Find dialog box, which is accessible through the View menu. Search for an event by one or more of the following: Type, Source, Category, Event ID, User, or Computer.

The settings in the Find dialog box are in effect during the current sessions. To save these settings, in Event Viewer, on the Options menu, click Save Settings On Exit.

Archiving Log Files

To archive an Event log, save it in one of three files formats:

- Log file format - Enables viewing of the archived log in Event Viewer
- Text file format - Enables presentation of the information in a text- oriented application such as a word processor
- Comma – Delimited text file format - enables management of the information with an application a such as a spreadsheet or a database

Archived Event logs can serve a variety of purposes, such as tracking and documenting system performance over time, or providing feedback to system and application developers. When Event logs are saved in any format other than their native.evt format, the hexadecimal data is lost

21.5 Short Summary

- ☞ Microsoft has provided an application programming interface (API) to enable developers to add objects to Performance Monitor that relate to their applications.
- ☞ When there is not enough physical memory installed, the paging file (PAGEFILE.SYS) is used as virtual memory.
- ☞ Performance Monitor is a great tool just to get a good idea of how your Windows NT computers are running and to analyze network segment traffic
- ☞ The counter Working set can tell you the amount of memory that the program has reserved for its use.
- ☞ An event is any potentially significant occurrence in the system or in an application.
- ☞ Windows NT records events in three kinds of log the System log, the Security log, and the application log.
- ☞ Archived Event logs can serve a variety of purposes, such as tracking and documenting system performance over time, or providing feedback to system and application developers
- ☞ Performance Monitor has four views : Chart, Alert, Log and Report
- ☞ All of the addresses intercepted by Network Monitor can be viewed through the Addresses command in the Capture menu.
- ☞ Network command in the Capture menu, track multiple network segments, which is attached to separate NICs.

21.6 Brain Storm

1. How can you audit the system performance?
2. What is the utility used to Monitor files operations in Windows NT?
3. Which view available through Performance Monitor should you use to create a baseline?
4. While using the Network Monitor, decide to implement a display filter to aid in your search for the NETBIOS Add Group Name command. What is the proper syntax of a properties-based display filter
 - a. NETBIOS : Command = 0x0 (Add Group Name)
 - b. NETBIOS <= Add Group Name Command
 - c. Add Group Name > NETBIOS
 - d. NETBIOS <- - > Add Group Name
5. Which application is best suited for detecting which computer on a local segment is causing the most network traffic?
6. What types of information can be tracked with Network Monitor?
7. Network Monitor's security feature of searching our other users of the Network Monitor Agent can detect what applications?
8. What is the use of Event Viewer?
9. What are the types of Events?
10. What are the types of Log?
11. How can you archive Log files?



Optimizing the Network Service Performance

Objectives

In this Lecture you will learn the following:

- Knowing the concept of Optimizing Windows NT File and Print servers
- Able to minimize the disk bottlenecks
- Knowing to eliminate the unneeded Network Protocols
- Able to limit the overcome network media
- Understand the concept of Optimizing Windows NT as an Application server

Coverage Plan

Lecture 22

- 22.1 Snap Shot
- 22.2 Optimizing NT File and Print Server
- 22.3 Optimizing NT as an Application Server
- 22.4 Short Summary
- 22.5 Brain Storm

22.1 Snap Shot

This lecture is divided into two principle sections. The first major section, "Optimizing Windows NT 4.0 File and Print Servers" illustrates how to use PerfMon to monitor specific characteristics of your network and computers and provides tips to obtain better network performance. The second section, "Optimizing Windows NT 4.0 as an Application Server", suggests techniques for maximizing the performance of Windows NT Server when running Windows NT services, such as SQL Server and Exchange Server.

22.2 Optimizing NT File and Print Server

Microsoft's developers made an appreciable improvement in Windows NT Server 4.0's file and print services compared with Windows NT 3.5x, especially for high-speed networks with 100Base T NICs. However, various hardware and software bottlenecks can still reduce the performance of Windows NT 4.0 file and print servers. The following sections describe the process of optimizing a Windows NT Server 4.0 used primarily for file and print sharing.

Minimizing Disk Bottlenecks

Fixed-disk drives contribute more than their share of performance problems. You use PerfMon's Physical Disk object to check for drive bottlenecks.

Unlike the other object in PerfMon, the disk subsystem must be initialized before disk activity can be monitored. To activate the monitoring of the disk subsystem, follow these steps:

1. From the Start menu, choose Programs and Command Prompt to open the Command Prompt window
2. On the command line, type `disdperf -y`.
3. Exit the command line by typing `exit` or by closing the Command Prompt window.
4. Shut down and restart Windows NT.

If performing these steps doesn't start the DiskPerf driver, the Physical Disk's counters don't work. Failure to start DiskPerf becomes evident when a Chart is selected for the physical disk and the graph doesn't indicate activity during disk operations.

One of the most obvious counters of the Physical Disk object is the %Disk Time counter, which is equivalent to monitoring the computer's disk activity indicator. When %Disk Time approaches 80 percent or higher, the server has a disk-usage problems and is said to be disk-bound.

Another useful counter is Disk Queue Length, which counts the number of processes waiting to use the physical disk. When more than two processes are regularly waiting to access the disk, the server is disk-bound.

Hardware Solutions for Reducing Disk Bottlenecks Buying SCSI-2 or, better yet, Ultra Wide SCST-3 fixed-disk drives with fast seek times is the most straightforward solution for increasing data throughput. Many disk drives are available with 4M/sec(4M per second) sustained data transfer rates or greater and 9ms seek times or less. For example 4.3 G Seagate ST15150W Wide SCSI-2 Barracuda drives can provide a sustained data rate in the range of 6M/sec, which is several times more data than you can transport over 10Base T media. A large read cache on high-speed drives also improves performance, and enabling a drive's write cache provides a performance boost, but at the expense of data security. The primary advantage of high-speed drives in relatively low-speed networks is that the system devotes less time to reading and writing data, which is only a part of the processor and network workload.

SCSI host controllers represent another potential hardware bottleneck. Attaching a high performance drive to a legacy SCSI controller (such as an 8-bit SCSI controller) doesn't make economic sense. PCI bus mastering controllers such as the Adaptec AHA-2940UW and AHA3940UW, provide synchronous data access at burst rates of up to 40M/sec and 80M/sec, respectively. You can attach up to 15 devices to a Wide SCSI-2 bus. Replacing Windows NT Server 4.0's software RAID implementation with a controller that implements hardware RAID also improves disk subsystem performance.

Some SCSI host adapters can perform asynchronous I/O, which allows drives to perform operations in parallel. If your host adapter can perform asynchronous I/O, you can use stripe sets to maximize the server's performance. A stripe set allows data to be distributed across several drives, making disk operations very fast because the physical drives work in parallel.

The cost of high-performance foxed-disk drives declined rapidly in 1996 and is expected to drop further in 1997 and 1998 as larger drives that use magneto-resistive heads and embedded servo positioning tracks become common. PCI bus-mastering host controllers range in street price from about \$350 to \$500, and low-cost hardware RAID controllers priced between \$750 and \$1000 became available in late 1996.

Software Solutions for reducing Disk Bottlenecks If your Windows NT server is used as a print server in addition to a file server, an option is to increase priority of the thread that handles file services and reduce the priority of the thread that handles print services. Roosting the priority of the file-server thread causes file requests to be handled more quickly, at the expense of print services.

To change the priority of the file-server thread with Registry Editor, follow these steps:

1. From the Start menu, choose Run to open the Run dialog.
2. In the text box, type regedt32 and click OK to launch Registry Editor(see figure 22.1)

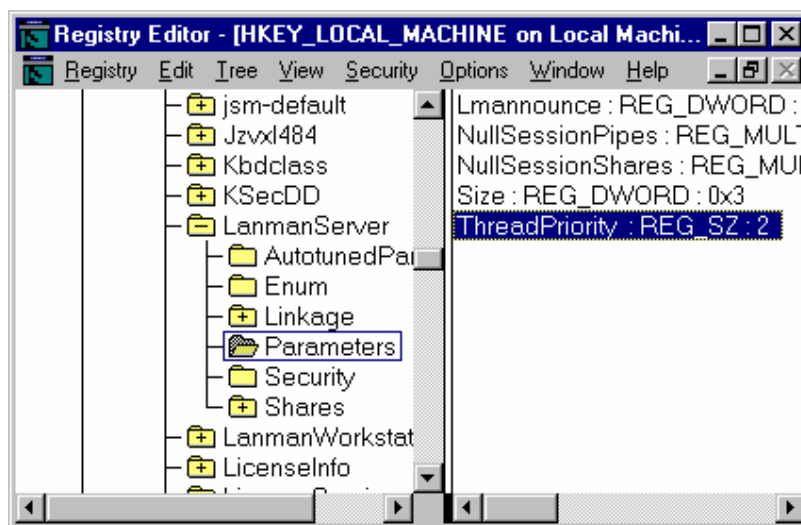


Figure 22.1 Using Registry Editor to change Windows NT system valves

3. Select the HKEY_LOCAL_MACHINE view.
4. Expand the SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters key.
5. Select the ThreadPriority entry.

NOTE : If your system doesn't have the Thread Priority entry, it can be added through the Edit menu. Set the initial value to 2

If your network consists of multiple Windows NT servers, you can use a technique called load balancing, which distributes the workload across multiple servers. Load balancing lets all the servers operate at a similar capacity, rather than heavily tax certain servers and under utilize other servers.

Before you load balance a system, evaluate each server to determine where the most activity occurs. Then the most heavily used files and folders can be replicated across other servers to balance the network load. Replicating is a method of duplicating the content of folders, files, or even entire disks onto another disk drive on another machine. (Although this technique is primarily used for providing an online path to information in case of a system crash of a primary serve, it can be used for load balancing also). Users that need to view information can be sent to the replicated data, thus reducing the traffic to the primary server.

Changing the Binding Order of Multiple Protocols

If you can't reduce the number of network protocols in use, you can alter the binding order to satisfy the most network requests on the first attempt. To view and alter the order of network bindings, follow these steps:

1. From the Start menu, choose settings and then Control Panel.
2. Double-click Control Panel's Network tool to open the Network property sheet.
3. Click the Bindings tab to display the Bindings properties page. Select the All Protocols entry from the combo box, if it's not already selected.
4. In the list box, you can change the order of the protocol bindings for each service. Double-click the entry for the service for which you want to change the binding priority to display the underlying bindings for the service
5. If most users use TCP/IP, for example, select this binding and click the Move Up button until TCP/IP is at the top of the list. (The Move Up and Move Down button are disabled until you select a protocol for a particular service)
6. Click OK to accept the changes or Cancel to abort the changes.
7. If you click OK and have made changes, Windows NT re-creates the bindings and asks that you shutdown and restart the server. On restart, the new binding priorities are effective.

Overcoming Network Media Limitations

You also can use Performance Monitor to view network characteristics that affect file server performance. Select the Server object and view the Bytes Total/sec counter : compare this number to the rated media speed of your network, such as 4 Mbps or 16 Mbps Token Ring or 10 Mbps Ethernet. If the

number of Bytes Total/sec is near the throughput limit of your network, your network is overworked. As an example, the maximum practical throughput of a conventional (Shared hub-based) 10Base T network is about 300,000 bytes per second : 100Base T delivers 3M/sec. In this case, you must separate the network into multiple segments to ease the burden. Fortunately, Windows NT Server 4.0 lets you install multiple network cards for segmenting. A more costly alternative is to replace the shared Ethernet system with switched Ethernet.

Some additional tips for improving the network performance of your Windows NT Server are as follows:

- ◆ If you're using 100Base T, change the server NIC(s) from the conventional PCI bus master variety to a model with an on-board processor, such as Intel's EtherExpress PRO/100 Smart Adapter, which has an i960 and 2M of RAM. The on-board processor reduces CPU loading.
- ◆ If your server uses WINS(Windows Internet Name Service) and TCP/IP, you can reduce the number of system broadcasts by binding the NetBIOS Interface to TCP/IP.
- ◆ If your network contains Windows for Workgroups (WfWg) clients and you use TCP/IP be sure to update the WfWg clients with the version of WfWg included on the Windows NT Server 4.0 CD-ROM. The new WfWg TCP/IP driver increases the server's performance.

Reducing File Fragmentation

It's a common misconception that Windows NT Server's NTFS eliminates problems with file fragmentation. Fragmentation – defined as files stored in multiple, non-contiguous clusters slows file read and writes operations, especially with large files.

Although Windows 95 includes a built-in de-fragmentation utility, Disk De-fragmenter, Windows NT 4.0 doesn't provide such an application. Commercial products are available to help correct and reduce the amount of disk fragmentation for your Windows NT systems. Symantec's Norton Utilities for Windows NT contains a utility similar to the Windows 95 Disk Defragmenter.

One unique product available is Diskeeper for Windows NT from Executive Software. Diskeeper is a Windows NT service that runs constantly in the background and controls disk fragmentation as it occurs. By having Diskeeper work in real time, you never have

to take the sever offline to perform defragmentation, and your drive is optimized for the fastest disk access possible.

22.3 Optimizing NT as an Application Server

Windows NT gained its initial reputations as a high-performance application server primarily by running client/server relational database management systems (RDBMSes), such as Microsoft SQL Server. The Microsoft BackOffice Server suite adds Exchange Server, Systems Management Server and System Network Architecture as server-based applications. Other firms, such as Oracle and IBM, offer server-based RDBMSes and messaging systems that run under Windows NT Server 4.0.

By running on the same system, several client/server applications that service large numbers or users can tax the capabilities of even the highest performance computer platforms. When dealing with application servers the most obvious place to look for bottlenecks is the CPU and memory systems.

Examining Memory Usage in an Application Server

One of the most effective methods for achieving better system performance is through the addition of more physical (DRAM or dynamic RAM) memory. Within reasonable limits, the performance of Windows NT Server 4.0's operating system and the server-based applications improves with the addition of more memory. Application servers generally start at 128M of RAM, and servers with 512M of RAM or more aren't uncommon. Just adding more RAM, however, isn't the entire solution to optimizing application server performance; you must correctly use the available memory.

Experience shows that no matter how much memory exists in a Windows NT Server system, the operating system and server-based applications always find ways to use additional RAM. Windows NT Server uses most memory as a disk cache.

You use the Memory object's counters in Performance Monitor to check memory usage. The number of Committed Bytes should be less than half-physical memory available, leaving the other half or more of RAM for disk caching. The number of Available Bytes should be at least 1M-preferably much more. If PerfMon displays less than the preceding values, you need more RAM.

Examining Virtual Memory Virtual memory supplements RAM with disk files that emulate RAM, a process called swapping or paging. Virtual memory is defined as the amount of physical memory (RAM) available in the system plus a preallocated amount of memory on the system's fixed disk. Windows NT

allocates the disk-based component of virtual memory in Pagefile.sys, which by default resides in the root folder of the disk drive on which you install Windows NT.

Pagefile.sys is unique because Windows NT creates the file as a series of contiguous clusters during installation. This structure lets the operating system issue a special set of disk I/O calls to read and write from Pagefile.sys, rather than use the conventional file system for disk I/O. The disk I/O calls to Pagefile.sys are much faster than Windows NT's normal disk operations but are much slower than memory I/O operations. To achieve maximum system performance, your objective is to minimize paging.

Virtual memory operations become very slow if the amount of virtual memory needed by Windows NT and running applications exceeds the amount of space allocated to Pagefile.sys. When this situation occurs, Windows NT dynamically expands the size of Pagefile.sys. The expansion process is slowed as Windows NT searches the disk for free disk space. Additional disk space is not in the contiguous-sector block of the original Pagefile.sys, a condition that slows read and writes operations.

Maximizing Network Throughput For application, network throughput is extremely important. The following steps help maximize network's throughput:

1. From the Start menu, choose Settings and Control Panel.
2. Double-click Control Panel's Network tool to open the Network property sheet.
3. Click the services tab to make the services page active, and select server from the Network Services list
4. Click the Properties button to open the Server dialog
5. Select the Maximize Throughput for Network Applications option.
6. If you don't have LAN Manager 2.x clients on the network, clear the default mark in the Make Browser Broadcasts to LAN Manager 2.x Clients check box.
7. Click OK to accept the changes, or click Cancel if you haven't made changes. If you have made changes, you must shut down and restart Windows NT to make the change effective.

To control the amount of virtual memory allocated to Pagefile.sys, follow these steps:

1. From the Start menu, choose Settings and then Control Panel.
2. Double-click Control Panel's System tool to open the system performance
3. Click the Performance tab to display the Performance properties page, and click the Virtual Memory button to open the Virtual Memory dialog
4. To change the amount of disk space consumed by Pagefile.sys, alter the setting dialog.
5. Click OK if you've made changes or Cancel if you haven't made changes. You down and restart the server if you make changes to Pagefile.sys

Viewing Virtual Memory with Performance Monitor The Commit Limit counter of PerfMon's Memory object shows the amount of memory in Pagefile.sys plus the amount of RAM that can be swapped to disk. In a system with 32M of RAM and a 57M Pagefile.sys file, the Commit Limit is 82.5M. When the number of Committed Bytes exceeds the Commit Limit, Windows NT tries to expand Pagefile.sys. To view the amount of memory that can't be paged from RAM, view the Pool Non-Paged Bytes counter.

There are four recommendations to reduce the amount of page swapping:

- ◆ Add more RAM to the system
- ◆ Remove unneeded services to lower operating system memory consumption.
- ◆ Remove unneeded device drivers.
- ◆ Install a faster fixed disk to reduce paging response time.

Two Windows NT Server 4.0 services consume large amounts of memory : DHCP (Dynamic Host Configuration Protocol) and WINS (Windows Internet name Service). If you need these services, you might find it a wiser choice to place DHCP and WINS on another server, such as a file server, rather than on the application server.

22.4 Short Summary

- Various hardware and software bottlenecks can still reduce the performance of Windows NT file and print Servers.
- Server threads run at foreground process priority by default.

- File server that's also a print server may suffer from server-thread starvation because the server threads are running at a lower priority than the print threads.
- Maximum value of Thread Priority is 31. Don't increase the Thread Priority value above 2; a value greater than 2 can cause other undesired system side effects
- Load balancing distributes the workload across multiple servers
- Multithreaded server applications take much better advantages of multiprocessor systems than file servers.
- Win32 synchronization object that ensures that only one thread can execute a particular block of code at a time.
- The amount of paging performed is directly related to how virtual memory is used.

22.5 Brain Storm

1. How you can reduce the file fragmentation?
2. Write the steps to minimize the disk bottlenecks?
3. Give a brief explanation on Optimizing Windows NT File and Print Servers?
4. Write the steps to change the binding order of Multiple protocols.
5. Give a notes on Optimizing Windows NT as an application server?
6. How you can examine the memory usage in a Application Server?



Lecture 23

Trouble Shooting

Objectives

In this Lecture you will learn the following:

- Able to detect and rectified the hardware problem
- Knowing about Network problems
- About the Primary troubleshooting tools
- Understanding the protocol analyzer

Coverage Plan

Lecture 23

- 23.1 Snap Shot
- 23.2 Hardware Problems
- 23.3 Boot Failures
- 23.4 Relating Network Protocols and Troubleshooting Issues
- 23.5 Network Problem
- 23.6 Using Protocol Analyzers
- 23.7 Windows NT Server 4.0's Primary Troubleshooting tools
- 23.8 Short Summary
- 23.9 Brain Storm

23.1 Snap Shot

Troubleshooting network problems is by no means an easy task. Network problems can arise in numerous areas; bad wiring causes problems at the physical layer, a failed network card results in trouble at the data-link layer and routing errors arise at the network layer. Worse yet, the network operating system or applications may cause problems at the transport session, resentation or application layer.

As this litany of potential woes indicates, it's important to have a good understand of the OSI model when you troubleshoot network problems. Finding a network problem often involves deciding which OSI layer is not the culprit. You need good troubleshooting tools and intimate familiarity with the networking protocols in use before you can be an effective network troubleshooter.

This lecture concentrates on solving the most common networking problems associated with Windows NT Server. Fortunately, Windows NT Server provides several tools (such as Control Panel) and utilities (such as Network Monitor) for configuring and troubleshooting networking problems.

23.2 Hardware Problems

The following section describes some of the hardware problems

SCSI problems

Some times problems that arise in installing Microsoft Windows NT Server have involved a SCSI controller. Under normal circumstances, a popular SCSI controller is identified properly. An improperly terminated SCSI chain, however, can create a condition that results in an inaccessible Boot Device error during the first reboot after the character portion of the installation program. It is very important to make sure that both ends of the SCSI chain are terminated. Even if DOS or another operating system is able to overlook certain conditions, Windows NT can be affected by even the smallest error.

Another small problem is during the SCSI detection routine. Most people don't think of looking on the list of SCSI devices for an ATAPI driver for their IDE based CD-ROM controller drive. Often, users are faced with a message asking for disk 4 because their CD-ROM controller was not detected. As more and more high speed CD-ROM drives are produced with an IDE interface, it is important to remember that if auto detection does not see your IDE controller,

you should press the S key to bring up the list of SCSI devices and choose the ATAPI 1.2 option to support your CD-ROM drive.

Hardware Driver Problems

Although the drivers that ship with Microsoft Windows NT Server usually are the latest at the time of shipping, manufacturers do update their drivers from time to time to reflect changes in the hardware that has been manufactured since the release date of Microsoft Windows NT Server. For this reason, whenever possible, find out whether there are updated drivers for your network interface card, and any other controllers, sound cards, or other devices. There is no guarantee that the drivers included with Microsoft Windows NT Server will work with your particular piece of hardware.

Also, although the Microsoft Windows NT Hardware Compatibility List might list your server, it also is possible that the manufacturer has shipped a model with slightly different hardware that might not be compatible with Windows NT. Most of the larger brand names perform testing to ensure that their hardware remains compatible, but the smaller players might not be as efficient. Be sure to get a commitment from your dealer or Value Added Reseller regarding Windows NT compatibility. This way, if it turns out that Windows NT cannot install on your server you have recourse and can get your server replaced with one that will give you a smooth installation.

Setup Disk Errors

Another problem that can come up is setup floppy disks going bad. This can happen with any floppy disk, so as Murphy's Law would have it, why shouldn't disk fail during Microsoft Windows NT Server installation?

You could use the winnt.exe option and install from MS-DOS (that is, if MS-DOS is already installed on the boot device,) but it's a good idea to have a good set of setup disks handy.

To create a new set of setup floppy disks, run the winnt.exe program from a DOS prompt. This program is located in the \i386 directory of the CD-ROM but add to that command line the /ox switch. This prompts you for three formatted 3.5inch disks, one at a time. When finished, you will have a brand new set of setup disks from which to work.

23.3 Boot Failures

Boot failures are problems during the startup of NT after a successful installation has completed. Some of the more common boot failures are presented in the sections that follow

NTLDR Error Message

If the NTLDR executable is missing or if a floppy is the A drive, then an error message is displayed as follows:

BOOT : Couldn't find NTLDR. Please insert another disk

If the NT DR is damaged or missing, you must use the ERD (Emergency Repair Disk) to repair or replace it. If there is a floppy disk in the drive, eject it, and continue the boot process

NTOSKRNL Missing Error Message

If the NTOSKRNL (NT Operating Kernel) is corrupt, missing, or the BOOT.INI points to the wrong partition, an error message like the following appears:

Windows NT could not start because the following
File is missing or corrupt:
\Winnt root\system32\ntoskrnl.exe
Please re-install a copy of the above file

This error can be resolved by repairing the NTOSKRNL file using the ERD repair process. Or if BOOT.INI is wrong, you can edit it to correct the problem

BOOT.INI Missing Error Message

If no BOOT.INI is present, the NTLDR will attempt to load NT from the default \ winnt directory of the current partition. If this fails, an error message similar to the following appears:

BOOT: Couldn't find NTLDR. Please insert another disk

To alleviate this problem, replace the BOOT.INI file from a backup or use the ERD to repair.

BOOTSECT.DOS Missing Error Message

If the BOOTSECT.DOS file is not present to boot to MS.DOS or another operating system (not NT), an error message appears as follows:

```
I/O Error accessing boot sector file  
multi (0) disk (0) rdisk (0) partition (1):\bootsect.dos
```

To indicate that the BOOT.INI file has been changed, the partition numbering has changed, or the partition is missing, inactive, or inaccessible. To attempt to repair or replace the BOOTSECT.DOS file, use the ERD procedure.

NTDETECT.COM Missing Error Message

If the NTDETECT.COM file is not present, the following error message appears:

```
NTDETECT V1.0 Checking Hardware...  
NTDETECT failed
```

This error must be repaired with the ERD repair process.

23.4 Relating Network Protocols and Troubleshooting Issues

Windows NT and its predecessor, LAN Manager, rely on Net BIOS protocols including server Messages Blocks (SMBs), named pipes and mail slots-for all native file and print services, plus many of Windows NT server's application services. Most Net BIOS - based protocol stacks, such as Windows NT's NETBEUI Frame (NBF), rely on broadcasting network packets to every client and server on the network.

Each Net BIOS device on the Windows NT network has a Net BIOS name, which uniquely identifies the device and contains information about the Net BIOS-related services the device provides. Broadcasting is means by which a Net BIOS device advertises its Net BIOS name and capabilities to other network devices. Broadcasting is also how a Net BIOS Device can locate another device or capability on network.

Broadcasting is deal for services that require real – time distribution to multiple clients, such as real time stock market data. Broadcasting is well suited for small networks because it reduces network response time. Reduced pocket overhead is responsible for NetBEUI's excellent performance on small networks. On larger networks, broadcasting creates a substantial amount of network traffic and adds to the difficulty of identifying a network device that's causing a problem .Much of your network troubleshooting time-if you have a large network is likely to be devoted to overcoming problems related to excessive broadcasts.

NetBEUI Broadcasting

All NetBEUI traffic is broadcast-based. If you install only the NetBEUI protocol on your Windows NT server, the server responds only to broadcast requests from other NetBEUI devices. NetBEUI isn't routable; if you have two Ethernet segments separated by a router, NetBEUI broadcasts aren't forwarded across the router. The only way to overcome this limitation is to enable bridging on the router, so that both appear as one physical segment. In this case, the router bridges the traffic between the two segments by using MAC (Media Access Control) layer addresses. Bridging eliminates the advantages of a segmented network. As you can imagine, maintaining a bridged network in a large LAN installation quickly becomes unmanageable, especially if you also must maintain routed protocols, such as TCP/IP, on the same segments.

If you must troubleshoot problems with a broadcast-based protocol, most of the tools for a directed protocol, such as TCP/IP, are ineffective. This is because many conventional network analysis tools rely on analyzing the interaction of source-destination network addresses to determine the flow of communication and where a connection is failing. Broadcast-based systems simply flood packets to every network device, whether or not the device is in the connection.

The best way to troubleshoot problems on broadcast network is to connect a protocol analyzer to the network, and then filter MAC addresses to determine who transmitted what to whom.

IPX/SPX

Windows NT 4.0's IPX support is similar to that for NetBEUI, except that IPX implements the OSI network layer, whereas NetBEUI doesn't. The addition of the network layer makes IPX routable; Windows NT's NW Link stack supports the IPX protocol with the Net BIOS session layer. This means that you can run your NT network by using only IPX for file, print and application services, if you choose.

In this case although the IPX/SPX protocol is routable, Net BIOS-over-IPX is broadcast-based. Net BIOS service announcement functions, such as browsing, aren't forwarded across subnets in a routed network. This means that devices running only NW link and separated by a router normally can't connect to each other's resources. You must have a mechanism on a dedicated routers, such as that provided by Cisco Systems' IPX Helper feature, to forward broadcasts to the desired destination.

TCP/IP

TCP/IP is a directed protocol that eliminates most of the traffic associated with NetBEUI and IPX protocols. A directed protocol usually involves point to point communication between two or more networked devices. Integration of NetBIOS and TCP/IP, which Microsoft calls NetBT (or, more commonly, NBT), follows two established Request For Comments (RFCs), defined by the Internet Engineering Task Force (IETF) as RFCs 1001 and 1002. From these RFCs, Microsoft built an entire suite of services to facilitate the use of TCP/IP with NetBIOS. The RFCs for NBT specify three TCP/IP service ports, which perform the following functions:

- NETBIOS Name uses UDP (User Datagram Protocol) port 137 for name-resolution requests.
- Net BIOS Data gram uses UDP port 138 for authentication, name registration, and browsing services.
- Net BIOS Session uses TCP (Transport Control Protocol) port 139 for server message blocks (SMBs) that file transfers and print jobs.

By using these three ports, Windows NT provides all its native services over TCP/IP. Windows NT also supports more traditional TCP/IP services, such as FTP and Telnet.

Windows Internet Name Service (WINS) use of TCP/IP requires traditional Net BIOS functions such as network browsing, name lookups, and user messaging to be mapped to the network addresses used by TCP/IP. Microsoft developed the Windows Internet Name Services (WINS) as its method for mapping IP addresses to Net BIOS machine names. The function of WINS is similar to the Domain Name Services (DNS) provided by most UNIX-based systems, but WINS provides additional Net BIOS functions for a given machine name to single IP address and to register Windows NT domain names to an IP address that represents a domain controller in the specified domain.

Machines running Windows NT server and Workstation with TCP/IP protocol stack use WINS to register their Net BIOS names and IP addresses. WINS registration occurs dynamically when the device starts up, or statically to guarantee that a certain machine name is registered to a certain IP address for a given user. For instance, the net send command queries WINS to find the IP addresses for the message destined for a specified user name.

After all servers and clients are registered with WINS, subsequent Net BIOS related operation such as browsing, messaging, authentication, and file and print services use the WINS database, located on a Windows NT server, to perform name resolution between Net BIOS names and IP addresses. Using the WINS database eliminates a need for broadcast name resolution thus decreasing network traffic and easing troubleshooting tasks.

The WINS database is a Jet (Access) database, Wins.mdb, and uses a Jet system database with the default name, Systems.mdb, both of which are located in \Winnt\system32\wins folder. Assuming that you've installed WINS, follow these steps to view the WINS database on your server:

1. From the start menu, choose Programs, Administrative Tools, and WINS Manager to open the WINS Manager window.
2. From the Mappings menu choose show Database to display the Show Database [Local] dialog

When you're troubleshooting WINS problem, it's important to keep in mind the different 16th Byte values (also called types). If the WINS database become corrupted, certain machine functions might be disabled because a 16th Byte entry for the service is missing or has the wrong value.

23.5 Network Problem

Choosing a Network that will not meet an organization's will lead directly to trouble. A common problem is choosing Peer - to - peer network when the situation calls for a server - base network

Trouble in a Peer-to-Peer Environment

A Peer-to-Peer or workgroup network may begin to exhibit problems with changes in the network site. These will not be hardware or software problem as much as logistical or operational problems. Indicators that a peer-to-peer network is not up to the task include:

- Difficulties caused by lack of centralized security
- Users turning off computers which are servers

Topology Problems

A network's design can cause problems if the design limits the network so it cannot perform in some environments

Bus Network

There are few situations that will cause a bus network to no longer be properly terminated and will usually take the network down. A cable on the network may

❖ Break

A break in the cable will cause both ends of the cable on either side of the break to no longer be Terminated. Signals will start to bounce and this will take the network down.

❖ Lose a connection

If a cable become loose and is disconnected, this will separate the computer from the network. It will also create an end that is not terminated, that will cause signals to bounce and the network to go down.

❖ Lose a Terminator

If a terminator is loose, it will create an end that is not terminated. Signals will start to bounce and the network will go down

Hub-Based Networks

While problems with hubs are infrequent they do occur. For example they can

❖ Drop a connection

If a computer becomes disconnected from the hub, the computer will be off the network but the rest of the network will continue to function normally

❖ Lose power

If an active hub loses power, the network will stop functioning

❖ Ring Networks

A ring network is usually very reliable, but problems can occur. For example a cable on the network may

❖ Break

If one of the cables in the ring breaks, that computer will be off the network but it will not affect the rest of the network

❖ Lose a Connection

If one of the cables in the ring becomes disconnected that computer will be off the network but it will not affect the rest of the network

23.6 Using Protocol Analyzers

Protocol analyzers are invaluable tools for discovering problems on a network or simply understanding the flow of packets between servers and clients. The advantage of dedicated protocol analyzer is its capability to capture problems from the OSI physical layer up to the application layer. If you have a Token Ring segment that's beaconing - that is generating an excessive number of beacon tokens because of some problem on the ring-you can see what device is responsible. If your network is experiencing intermittent loss of connectivity, a protocol analyzer might show you that a specific server or client workstation is generating a large number of malformed packets, in which case you remove and repair the offending device. It's particularly important that the hardware or software protocol analyzers you use recognize the unique nature of Windows NT specific operation such as DHCP, WINS and SMBs.

Hardware Protocol Analyzers

The best known hardware-based protocol analyzer is Network General's Expert Sniffer. The Sniffer usually comes bundled with a portable PC and includes a specialized network interface card that's designed to capture and decode packets. Network General also supplies PCMCIA cards and software for use in specific notebook PCs, which let you build your own Sniffer. Network General also produces the Distributed Sniffer, a specialized box that lets you plug into

and capture packets on multiple segments simultaneously. Devices such as the Expert Sniffer not only let you capture and decode packets, but also provide help with diagnosing problems by analyzing the data and suggesting possible causes for detected problems.

Software Protocol Analyzers

Software-only protocol analyzers cost less than hardware devices but provide less functionality. Examples of popular software-based analyzers are Novell's LANalyzer for Windows and Microsoft Network Monitor, which is part of Microsoft Systems Management Server (SMS) and is now included with Windows NT Server. Both applications provide the capability to decode network packets and support various types of pre and post packet filtering for several protocols. Network Monitor is ideal for troubleshooting Windows NT networking problems because it's designed to identify the NetBIOS elements specific to Windows NT networks.

Protocol Analyzer Connection

All protocol analyzers, whether they're hardware or software based, must be connected to each segment you want to monitor, you need an analyzer agent on each segment of a routed network to capture packets for analysis. The agent is necessary because routers don't forward some packets that may be of interest, such as broadcast packets.

23.7 Windows NT Server 4.0's Primary Troubleshooting Tools

Windows NT provides a variety of built-in tools for troubleshooting network problems. Using Windows NT tools is important because traditional network troubleshooting tools don't always support the NetBIOS implementation of network protocols that Windows NT requires.

Using Event Viewer

Make Event Viewer's System log your first step in the diagnosis process. From the Start menu choose Programs, Administrative Tools(Common) and Event Viewer. Event Viewer's three classes of errors are as follows:

- *Error* also called a stop event, which indicates a significant problem, such as failure of the server’s network card or the inability of a network service to load. A stop event isn’t the same as a stop error; stop errors are fatal, low-level errors.
- *Warning* which indicates a problem of less severity than a stop event, such as failure of a connection to a backup domain controller or a master browser
- *Information* which describes successful completion of important events, such as starting SQL Server or Exchange Server.

Look for network-related errors, including BROWSER, WINS and NeBT errors and warnings. Figure 23.1 shows Event Viewer displaying stop events for Remote Access Service and dependent services caused by a modem hardware failure and a warning from the BROWSER service. Double-clicking the BROWSER warning opens the Event Detail dialog shown in the figure 23.2, which indicates that the NTWS1 primary domain controller can’t communicate with the NTWS1 backup domain controller, which also is a browse master.

Date	Time	Source	Category	Event
24.7.2001	7:36:48	NNTPSVC	None	116
24.7.2001	7:36:48	Service Control Mar	None	7023
24.7.2001	7:36:48	MSFTPSVC	None	116
24.7.2001	7:36:37	Server	None	2511
24.7.2001	7:36:36	Server	None	2506
24.7.2001	7:36:35	Service Control Mar	None	7000
24.7.2001	7:36:21	EventLog	None	6005
24.7.2001	7:36:21	EventLog	None	6009
24.7.2001	7:36:30	qic117	None	119
23.7.2001	21:01:48	EventLog	None	6006
23.7.2001	21:01:48	BROWSER	None	8033
23.7.2001	21:01:46	BROWSER	None	8033
23.7.2001	21:01:46	BROWSER	None	8033
23.7.2001	20:32:36	Srv	None	2013
23.7.2001	20:28:12	Service Control Mar	None	7026
23.7.2001	20:27:48	Service Control Mar	None	7023

Figure 23.1 Event Viewer’s system log showing stop events for RAS and the browse master failure warning

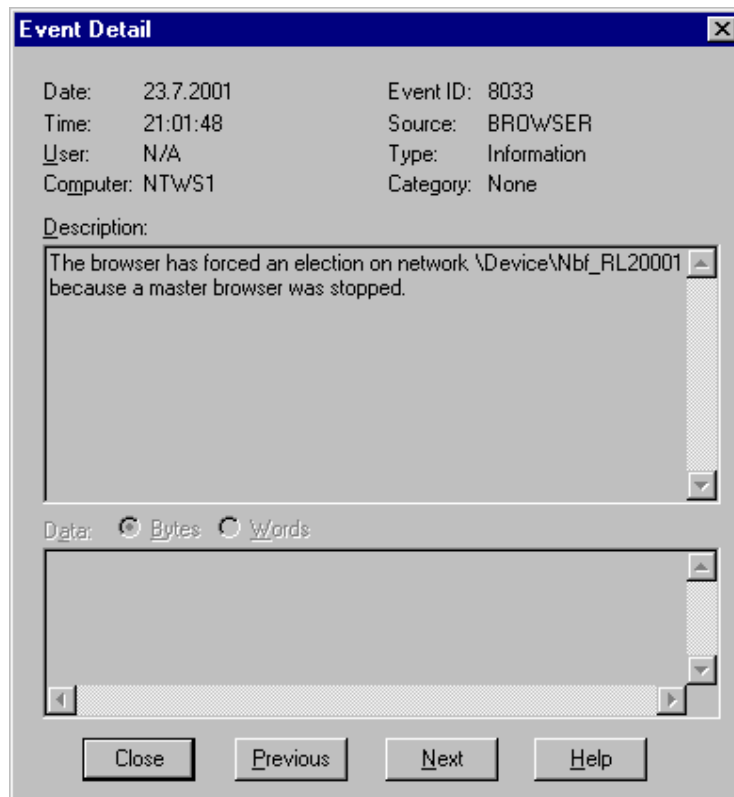


Figure 23.2 The Event Detail dialog for the BROWSER warning of figure 23.1

Using Network Monitor

The Network Monitor tools is very handy for troubleshooting NT related network problems.

Installing NetMon and its Agent Service You have two options when installing NetMon components on NT Server: you can install just the Network Monitor agent, or install both the NetMon GUI tool and the agent. You choose the agent only install if you have the SMS version of NetMon and want to enable monitoring across a routed network. In this case, the agent-only option is installed on one NT Workstation or Server per segment.

To install the Network Monitor Tools and the Agent Service follow these steps:

1. From Control Panel, double click Network to open the Network property sheet and then click the Service tab.
2. Click Add to open the select network service dialog

3. In the Network Service list, select network monitor tools and agent; then click Have Disk to open the Insert disk dialog.
4. Enter the path to your NT server distribution CD-ROM and the appropriate subfolder for your processor. For instance, if your CD-ROM drive is E on an Intel-based server, type e:\i386 and click OK
5. After finishing copying files, you return to the Network property sheet. Click close to reconfigure network bindings and request a system restart. Restart the server for the changes to take effect.
6. By default, the Network Monitor Agent Services is installed for manual startup. To start the Agent Service, type net start "Network Monitor Agent" at the command prompt or from the Control Panel's Services tool, select Network Monitor Agent and click Start.

Putting Windows NT Network Monitor to use To use NetMon follow these steps:

1. From the Start Menu, choose Programs, Administrative Tools and Network Monitor to launch NetMon
2. Choose Networks from the Capture menu to open the select capture network dialog. You must connect to the Network Monitor agent bound to your server NIC before using NetMon. If you have more than one NIC in the server you can bind to one or the other to trace packets.
3. Select the desired adapter to connect, if you have more than one NIC, and click OK to close the dialog.
4. Choose Start from the Capture menu to open the Capture window and begin capturing packets

By default the capture window is broken into four panes: Total Statistics, Session Statistics, Station Statistics and Graph. You can toggle which panes you want to see by clicking the four pane toggle buttons of the toolbar. Alternatively, you can mark or unmark the appropriate checked item of the Window menu.

5. After you finish capturing, from the Capture menu choose Stop or Stop and View. If you choose Stop and View the Capture summary window appears.

6. Choose Filter from the capture summary windows display menu to open the Display Filter dialog which lets you create filters for the captured packets. You can filter by protocol, by station address or by protocol property. If you want to set up a filter so that only IP packets appear, click Disable All in the Expression dialog; then from the Disable Protocol list, scroll down to select IP. Click enable to display only IP packets. Click OK twice to close the dialogs and engage the filter.
7. Double -click a packet item in NetMon's list to display additional information about the packet. The Capture Summary window opens with three panes that provide summary, detail and hexadecimal information about the selected packet.
8. Click the middle pane to examine the packet at each protocol layer. Each OSI layer is represented by an entry including Physical, MAC(Ethernet,Token Ring and so on), Network(IP,IPX,AppleTalk,and so on), Transport (TCP) and Session (NBT).

A plus symbol to the left of each layer item indicates that more detail is available under the heading. Double-click one such entry to display detailed information for the portion of the frame.

Using Performance Monitor

Windows NT's Performance Monitor (PerMon) is a valuable tool for monitoring the operation of NT servers and workstations, including their network components. PerMon lets you monitor network interfaces on a server to determine bandwidth usage, rates of errors and broadcasts and protocol-specific counters.

Installing the SMTP Service You need to install the SMTP service to obtain all the available TCP/IP or IPX statistics for your network interface. Without SNMP, PerMon can't see some of the network objects. To install the SNMP service follow these steps:

1. From the Control Panel, double-click the Network tool to open the Network properly sheet and then click the Service tab.
2. Click Add to open the Select Network Service dialog.

3. In the Network Service list, select the SNMP Service and then click Have Disk to open the Insert Disk dialog
4. Enter the path to your NT server distribution CD-ROM and the appropriate subfolder for your processor. For instance, if your CD-ROM drive is E on an Intel-based server, type e:\i386 and click OK to copy the SNMP service files.
5. The Microsoft SNMP Properties sheet appears, letting you customize the SNMP agent on your system
6. After you complete the setup of the SNMP agent for your server, click OK to close the Microsoft SNMP Properties sheet and click Close in the Network property sheet. Restart your system for the changes to take effect.

Using PerMon with TCP/IP Networks After installing the SNMP service on your NT server, you can use PerMon to gather IP or IPX related statistics about your system. The most useful application of PerMon as a network troubleshooting tool is tracking protocol-related information over time with PerMon's logging function.

PerMon also includes counters for ICMP, TCP and UDP objects that provides similar information to IP. Tracking these counters over time provides valuable information about network performance on your server, especially when you must troubleshoot TCP/IP or IPX network problems.

Using Network -Related Command -Line Tools

Windows NT provide many useful command-line tools for troubleshooting network problems. The following sections describe the most important tools included with server. You can use the command-line tools that come with NT help solve many basic networking problems. Most of the tools are for TCP and are likely to be familiar to UNIX users.

Address Resolution Protocol (ARP) The arp command lets you view the current contents of the ARP cache on a server or workstation. Arp-a displays the contents of the ARP cache; arp-d and arp-s let you manually remove and add entries to the ARP cache. With ping you can use arp to determine whether a device is communicating on the network. If you ping a device in question, you should see a corresponding entry in the ARP cache of either device you're pinging, or that of the default gateway if the device isn't on your local subnet.

Hostname The hostname command returns the name of the system on which the command is executed. The name returned is the name specified in the DNS setup section of the TCP/IP configuration of the system, rather than the NetBIOS name.

Ipconfig the ipconfig command returns all the current TCP/IP,DNS and WINS information for the system. Use ipconfig/all to display all the information and just ipconfig to obtain abbreviated information.

The ipconfig command is also a quick way to get the MAC address of the NIC installed in the device. If you're using DHCP, you can also use ipconfig to renew a DHCP address reservation or release an address on a DHCP client by using the /renew and /release parameters.

Nbtstat The nbtstat command is one of the most useful networking tools because it provides various information about NetBIOS names and their addresses. For example if you know the NetBIOS name of the workstation and want to know its IP address, follow these steps:

1. Type **net view \\ machine name** at the command prompt, where machine name is the NetBIOS name of the device. You receive a list of shares available on that machine or the message *There are no entries in the list*
2. Type **nbtstat - c** to display the name and the IP address of the machine specified in step 1. You don't need to specify the machine name because the result of the preceding name resolution is cached in the NetBIOS Name Cache, which you can view with the -c parameter.

You also use the nbtstat - A ip-address command to determine what machine is registered to a given IP address. Note that this command requires an uppercase A parameter. When you issue the command, the server or workstation sends a name request to the IP address of the primary WINS server specified in the issuing device's TCP/IP WINS configuration page. The returned information is the contents of the WINS database for ip-address. The command is useful if you're trying to troubleshoot WINS problems.

23.8 Short Summary

- ⊗ A Peer-to-Peer or workgroup network may begin to exhibit problems with changes in the network site.
- ⊗ A ring network problems can occur when the network cable may Break or Lose a connection.

- ✗ A Hub based Network problems may occur when the network drop a connection or lose power.
- ✗ In event viewer all the logs collect the same information about each event: date, time, source, category, event, user ID.
- ✗ Most system errors, including stop errors that result in the blue screen, are recorded in the System log. This allows you to review the time and circumstances around a system failure.
- ✗ If the NTLDR executable is missing then the error message will be displayed.
- ✗ If the NTOSKRNL (NT Operating Kernel) is corrupt, missing, or the BOOT.INI points to the wrong partition

23.9 Brain Storm

1. In which situations network problems are occurring in the following networks:
 - a. Bus network, b. Hub based network, c. Ring based network
2. Explain the tools that can be used to repair and correct operational difficulties?
3. What is the Meaning of LKGC?
4. Which of the following can be corrected by the repair process using the three installation disks and a recent Emergency Repair Disk
 - a. Boot Sector corruption
 - b. Unable to locate Master Boot Record
 - c. NTLDR not found
 - d. Corrupt NTOSKRNL
5. Which two parameter switches are present by default on the VGA Mode selection ARC name line in the BOOT.INI file?
6. What are all the error messages displayed when the system gets boot failure?
7. What is NTLDR?
8. After installing a new SCSI driver, NT will not successfully boot. No other changes have been made to the system. What is the easiest way to return the system to a state where it will boot properly?
9. How can you rectify Hardware Driver Problems?
10. How do you solve Setup Disk Errors?

☞ Best of Luck ☞

Operating Systems with Windows NT/2000

Lecture 1 Introduction to Networks

Introduction to Networks - About PC Network - Concept of Networking - Benefits Of Networking - Types Of Networking - Classification Of Networking

Lecture 2 Introduction to NT

Introduction to NT - Networking with Windows NT Server 4.0 - Windows NT Server 4.0, the Internet and Intranets - What New in Windows NT Server 4.0 - Coming Attraction in Windows NT 5.0

Lecture 3 Understanding the Windows NT Operating System

Understanding the Windows NT Operating system - Windows NT Operating System Features - Windows NT System Architecture

Lecture 4 NT Environmental Subsystems

NT Environmental Subsystem -Client and Protected Subsystem Servers

Lecture 5 Choosing Network Protocol

Choosing Networking Protocol - Understand the OSI Seven-Layer Model - Comparing Windows NT and OSI Network Layers - Networking with Windows NT's Protocols - Supporting a Variety of PC Clients

Lecture 6 Network Topologies & Architecture

Network Topologies & Architecture - Access Methods - Network Topologies - Network Architecture

Lecture 7 Transmission Media

Transmission Media - Cable Media - Wireless Media - Comparisons of different Wireless Media

Lecture 8 Network Adapter Card

Network Adapter card - Working of a Network Adapter Card - Network Adapter Card Compatibility - Configuring Network Adapter Cards

Lecture 9 Connectivity Devices and Transfer Mechanism

Connectivity Device & Transfer Media - Addressing - Modems -
Repeaters - Hubs - Bridges - Routing - Gateways

Lecture 10 File Systems

File System - NTFS File System - Understanding NTFS Permission -
Compressing NTFS Files and Folders

Lecture 11 Computer Security

Computer Security - Security & Windows NT - C2 Security - Requirements
For C2 Security

Lecture 12 Windows NT Server Installation

Windows NT Server Installation - Installing Windows NT Server - System
Requirements - Compatibility Issues - Types of Installation - Starting the
Basic Installation - Repairing the Windows NT Server Operating System
Installation

Lecture 13 Setting up RAID

Setting up RAID - Understanding RAID Levels - Creating Windows NT
Server Stripe and Mirror sets - Recovering a Software RAID 1 or RAID 5
Set

Lecture 14 Installing File Backup Systems

Installing File Backup Systems - Backup Types - Developing a Backup
Strategy - Choosing Backup Hardware - Windows NT Server 4.0 Backup
Application

Lecture 15 Windows NT Registry

Windows NT Registry - Registry Basics - Configuration settings in
Registry - Registry's Organization - Registry Editor - Important Hives &
Keys - Inspecting Another Computer's Registry - Maintaining Registry
Security

Lecture 16 Using TCP/IP, WINS and DHCP

Using TCP/IP ,DHCP and WINS - Role of TCP/IP - Installing & Configuration TCP/IP - Implementing DHCP - Implementing WINS

Lecture 17 Working with Domains

Working with Domain - Win NT Domain Models - Domain Architecture & Security - Implementing Domains and Trusts between Domains

Lecture 18 Managing User Account

Managing User Account - User Account - Working with User Manger for Domains - Managing User accounts & their Properties - Administering the Domain Account Policy

Lecture 19 Managing Group Accounts

Managing Group Account - Group Account - Managing User Groups - Using Group Management Wizard - Managing User Rights Policy

Lecture 20 Sharing and Securing Network Resources

Sharing and Securing Network Resources - Sharing & Securing Folders and Files - Replication Folders - Sharing and Securing Network Printers

Lecture 21 Monitoring the Network

Monitoring the Network - Performance Monitor - Network Monitor - Event Viewer and Log Files

Lecture 22 Optimizing the Network Server Performance

Optimizing the Network Server Performance - Optimizing NT File and Print Server - Optimizing NT as an Application Server

Lecture 23 Troubleshooting

Troubleshooting - Hardware Problems - Boot Failure - Relating Network Protocols and Troubleshooting Issues - Network Problem - Using Protocol Analyzers - Windows NT Server 4.0's Primary Troubleshooting tools

∞ End ∞