Network Naming And Security

Terms you'll need to understand:

- √ Network names
- √ Accounts
- √ NetBIOS computer names
- √ Windows Internet Name Service (WINS)
- √ Domain Name Service (DNS)
- √ LMHOSTS file
- √ HOST file
- √ Security policy
- √ Virus protection
- √ Disaster recovery
- √ Backup
- √ Uninterruptible Power Supply (UPS)
- √ Redundant Array of Inexpensive Disks (RAID)

Techniques you'll need to master:

- √ Planning and implementing network security
- √ Enabling network auditing for security purposes
- √ Implementing a disaster recovery plan

As computers become increasingly capable of communicating with one another, and individuals become more knowledgeable about operating these computers, security becomes an issue of paramount importance. This chapter covers computer security relating to the following areas: network naming schemes, planning for network security, security models, additional security considerations, and disaster recovery.

Network Naming Schemes

Computers identify themselves to each other using various naming schemes. Each computer in a networked environment must be assigned a name so it can communicate with other computers on the network. Additionally, users on the network need names, as do the shared resources on the network. These network names can be divided into the following categories:

- Accounts
- Computers
- Resources

Accounts

An account is a compilation of all the information that pertains to a particular user or group on the network—normally consisting of a user name, password, rights and permissions, and group memberships. Accounts are created by the administrator, and are required to operate in most secure systems. It is in the best interest of network security for a user not give other users his or her password, or any other account information that may be used to compromise sensitive data or operations.

Computer Names

One individual can have many titles simultaneously—a man may be called a father, boss, neighbor, or son, depending on who is addressing him. In much the same way, each computer in a network can have multiple names, depending on what process, protocol, or device is communicating with it at the moment. Because not all entities use or understand the same naming schemes, you need a system that allows for translation (resolution) of one naming/addressing type to another.

For instance, a data frame received at a computer's physical layer contains information specifying that it is addressed to the computer's NIC address (or MAC address). The MAC address consists of hexadecimal code stored

on a ROM chip on the computer's NIC; and at the physical level, this computer is known by this code. Numbering systems like this are easy for computers to understand, but they are very difficult for humans to use and remember. However, a user-friendly name like "Print Server2" would be easy for people to use.

In a similar manner, remote systems using different equipment or protocols may have difficulty deciphering your naming and addressing scheme. To address this problem, various name resolution standards have been created to help transparently resolve the different names to one another, allowing internetwork communication to occur.

NetBIOS Computer Names

Every Microsoft computer uses a computer name, also known as the NetBIOS computer name, of up to 15 characters (see Figure 9.1). These names are part of the OSI application layer interface to the network protocol stack on all networked Microsoft computers. Networks that use the TCP/IP protocol must resolve (or translate) the computer name into an IP address before networked communication can occur. The Windows Internet Name Service (WINS) and/or the Domain Name Service (DNS) can be used to resolve computer names to their corresponding IP addresses.

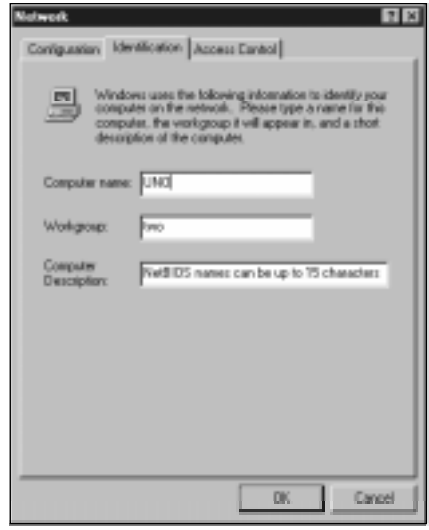

Figure 9.1 NetBIOS names can be set up to 15 characters in Microsoft computers. The name of this Windows 95 computer is "UNO."

Windows Internet Name Service

WINS is a service that resolves NetBIOS computer names to IP addresses. This service runs on a Windows NT server on the network, and dynamically resolves NetBIOS computer names to IP addresses so that computers can communicate with each other. WINS is a client/server service whereby a client registers its computer name with the WINS server during its boot process. When a WINS client needs to locate a computer on the network, it can query the WINS server to obtain a computer name to IP address mapping (translation) for the computer with which it is going to communicate (see Figure 9.2).

Domain Name Service

DNS is similar to WINS in that it resolves computer names to IP addresses. However, in DNS, these names are called "host names" or fully qualified domain names (FQDNs). In general, NetBIOS computer names consist of a single part. In contrast, TCP/IP components rely on a naming convention know as the Domain Name System (DNS). A FQDN is a hierachical naming convention using the format hostname.domainname, such as microsoft.com, spca.org and utexas.edu, where the domain name is an indication of the type of organization. Windows NT combines the NetBIOS computer name with the DNS domain name to form the FQDN.

Figure 9.2 The WINS server can resolve NetBIOS names to IP addresses.

DNS also differs from WINS in that DNS is a static service, meaning that someone has to manually enter the names and IP addresses before the DNS server can resolve them. Figure 9.3 illustrates how the DNS server handles a request from the DNS client. The DNS server also can query other DNS servers to retrieve partial resolutions for the computer name. For instance, one DNS server may resolve the "microsoft.com" part of the name, while another resolves the "uno."

In addition to WINS and DNS, LMHOSTS and HOST files can be used to resolve names to IP addresses.

LMHOSTS And HOST Files

LMHOSTS and HOST files are normally maintained locally on the client computer. These files must be manually created and updated and placed in the appropriate directory (for instance, WINDOWS\ETC in a Windows 95 computer or WINNT\SYSTEM32\DRIVERS\ETC in a Windows NT computer). The LMHOSTS file is normally responsible for resolving NetBIOS names to IP addresses; the HOST file resolves host names and/or FQDN names to IP addresses.

Resources

Each resource is identified by a name. These names can be as vague or specific as the individual sharing the resource desires. "Printer connected to LPT 1" and "High-speed color thermal imaging printer on second floor,"

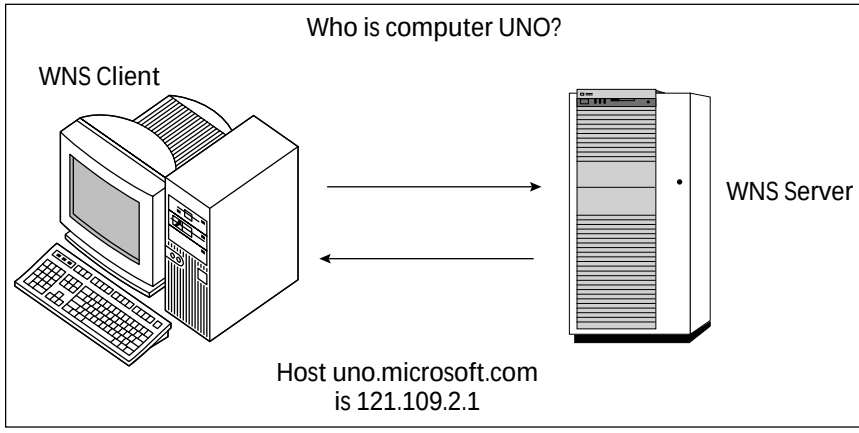

Figure 9.3 This graphic illustrates how DNS can resolve fully qualified domain names.

are both possible names for the same resource. Optimally, resource names should be chosen to make it easy for a user to tell what object (printer, shared file, or directory) is being accessed. Figure 9.4 illustrates the different types of resources you may see when browsing the network.

Planning For Network Security

If you don't take the time to plan your network, you can spend a huge amount of time, effort, and money correcting errors and redesigning later should you find that it isn't performing as you expected. This becomes an even larger issue in the area of security. You must create security policies and plans based on your organizational requirements. You should ensure that your network plan does the following:

- Addresses your company's needs
- Establishes policies that support those needs
- Provides the necessary amount of physical and logical security

Understanding Requirements

Your organization is different from every other organization in the world. It has a unique personality, style, and sense of direction. Your security plan must fit within your organization's style. Although you may think the CEO and everyone else should know your security ideas, remember that those ideas must support the direction of the organization. For instance, requiring diskless workstations in a public lending library probably won't fit in with the head librarian's vision of sharing information.

Look at your organization's vision. If it doesn't have one, look at other organizations in similar fields. Do you want an Internet presence? Do you

Figure 9.4 Shared resources are usually printers, drives, or directories made available by other computers on the network.

want one in the future? Is your data so sensitive that it must be kept absolutely private? Do you want to be able to remotely administer your site? You and your company's leadership must decide the depth of security for your network. Then, make the system as easy to use as possible, without violating your original security goals. Security threats to your network include:

- Unintentional damage
- Intentional damage
- Unauthorized access
- Electronic tampering
- Theft

Your security policy should address each of these threats in detail.

Setting A Security Policy

You are responsible for your organization's security policy. This policy must describe the who, what, when, where, and possibly the why of how your organization implements security. It is probably not a good idea, however, to describe exactly how your security is implemented.

You should maintain a network log, separate from your published security policy, which describes all actions you have taken in setting up and supporting your network. This is where your fellow administrators can look to see what has been implemented.

Many users feel that security policies are hindrances to getting work done, and will attempt to "get around" them. Don't immediately discard their ideas. Instead, review them, and after weighing the benefit of accomplishing work more easily against the benefit of additional security, you can decide whether to change the policy or leave it in place. Explain to the users why the system was or was not changed, and thank them for their suggestions.

Users can accidentally or unintentionally damage data, or the network itself. Policies must allow users to easily and efficiently accomplish their work, without allowing them to cause any damage.

Physical And Logical Security

No matter how well you have your data protected electronically, it doesn't do any good if your computers can be carried out the front door of your

building. You are responsible for all aspects of securing your network, including physically securing it.

In a peer-to-peer network that uses shared resources, each user is responsible for securing his or her computer. Your policy may be as simple as having users shut down their computers every night and locking their office doors. Some users may prefer to leave their computers on at night so they can obtain files from their home—you must decide if your security policy will allow for this type of situation. Again, you must balance the needs of the users with the need for secure network management policies.

In a client-server network, each user is allowed specific rights based on his or her needs and role within the organization. Like users in peer-to-peer networks, users in client-server networks need to take a degree of responsibility in keeping their personal workstations secure.

Servers

As the administrator, you must prevent internal and external sources from compromising your servers. A well-intentioned user can cause irreparable data loss by tinkering in the server closet. You should limit physical access to your servers, as well as limit the users and groups that can log on locally to the server.

Routers

Routers, much like servers, should be allowed limited access. A user might decide that because the network was acting oddly, maybe turning the router off and back on could fix the problem. Depending on the make and model of the router, it may actually be okay to do this. Or it may delete the routing tables, and cause the router to "forget" what it is. Also, if your organization is large enough, a user may add his or her own network segments to your routers, possibly without your knowledge. This could be the first step in a major compromise of your network.

Cables

If your data is sensitive enough, you may have to limit the amount of electronic signals your cables emit, or ensure your topology doesn't radiate these signals. Also, it is possible for someone to tap directly into copper cable and thereby steal data. (During the mid-1970s, at the height of the Cold War, the United States government developed a wireless technology that can eavesdrop on the electromagnetic signals emanating from computers and network cabling.)

You should limit access to cable runs that carry sensitive data. If at all possible, use fiber optic cable, which is much more difficult to tap and does not emit electrical signals. Also, use the building structure to protect your cables. For instance, you may be able to run your cabling inside the walls of the building.

Security Models

Once you have physically secured your network, you must decide how you will secure it electronically. There are two basic types of account security that you can implement, share-level and user-level security, which we introduced in Chapter 6.

Share-Level Security

Share-level security is a security technique where each resource owner shares a resource under his or her control and creates a password to control access to this resource. An example of this would be a user with a color printer connected to his computer. He could share the printer over the network, and protect its use by placing a password on the use of the printer. Users who know the password can use the printer. Windows 3.11 and Windows 95 are examples of operating systems that use this technique. Both of these systems allow Read Only, Full, and Depends On Password access to the resource (see Figure 9.5), each of which is explained in the following:

- Read Only Users are only allowed to view files in the shared directory. They may not change, delete, or add files to the directory. They may copy files to their computers, and print (or execute) them.
- Full Users may create, read, write, delete, and change files in the shared directory.
- Depends On Password Shares using this feature allow users either Read Only rights, or Full rights, depending on the password the user provides.

User-Level Security

User-level security is a security technique where each user is assigned a unique user name and password. When users attempt to enter the network, they are prompted to provide their user name and password, which are compared to a security database on a remote server. This process is known as authentication. If the user name and password are correct, the server logs

Figure 9.5 Share-level security provides password protection capabilities for shared resources.

the users onto the network. Rights and privileges are assigned based on the user ID and/or groups to which the user may belong.

In Microsoft NT, this technique is used to assign permissions to. Although users may have access to parts of the network, they may or may not have permission to access certain resources, or their type of permission may limit their level of interaction with these resources. Those types of permissions are as follows:

- Read Similar to the Read Only permission in Windows 3.11 or Windows 95, the Read permission allows users to view files in the shared directory. They may not change, delete, or add files to the directory. They may copy files to their computers and print (but not execute) them.
- Execute The Execute permission allows users to run files in the shared directory.
- Write The Write permission allows users to create, read, write, and change files in the shared directory. Users may not execute or delete files.
- Delete The Delete permission gives users the ability to delete files in the shared directory.

- Full Control Full Control access allows users to read, execute, create, write, change, and delete files in the shared directory.
- No Access Users are not allowed to access this directory if the No Access permission is set. When combined with other permissions, the No Access permission takes precedence. That is, if a user has Full Control and No Access to a particular directory, the user has No Access to the directory.

Account Management

To keep track of users and permissions, the administrator logically divides users and resources into manageable groups, and assigns permissions to each group.

For example, an organization consisting of Management, Engineering, Marketing, Accounting, and Human Resources could easily be grouped according to each specialty. In this instance, the administrator would create account groups (called global groups in Window NT) named: Management, Engineering, Marketing, Accounting, and Human Resources. Then, the administrator would place the users in each discipline into the appropriate group. As each group will need access to different resources, the administrator would create resource groups (called local groups in Windows NT) with the permission to access the following devices: High-speed printer, Color printer, Printer, CAD printer, After-hours printer, and Payroll printer. The administrator would then create account groups: High-speed printer with Management and Marketing; Color printer with Marketing; Printer with Accounting, Engineering, and Human Resources; CAD printer with Engineering; After-hours printer with Everyone; and Payroll printer with Accounting. (see Figure 9.6).

Additional Security Considerations

The security measures that we've already mentioned are essential to your network's security. Other techniques you may consider implementing are auditing, diskless workstations, encryption, and virus protection.

Auditing

Auditing is a way to monitor network events and user actions. Primarily, auditing provides a log trail that shows who did what, and when, on the

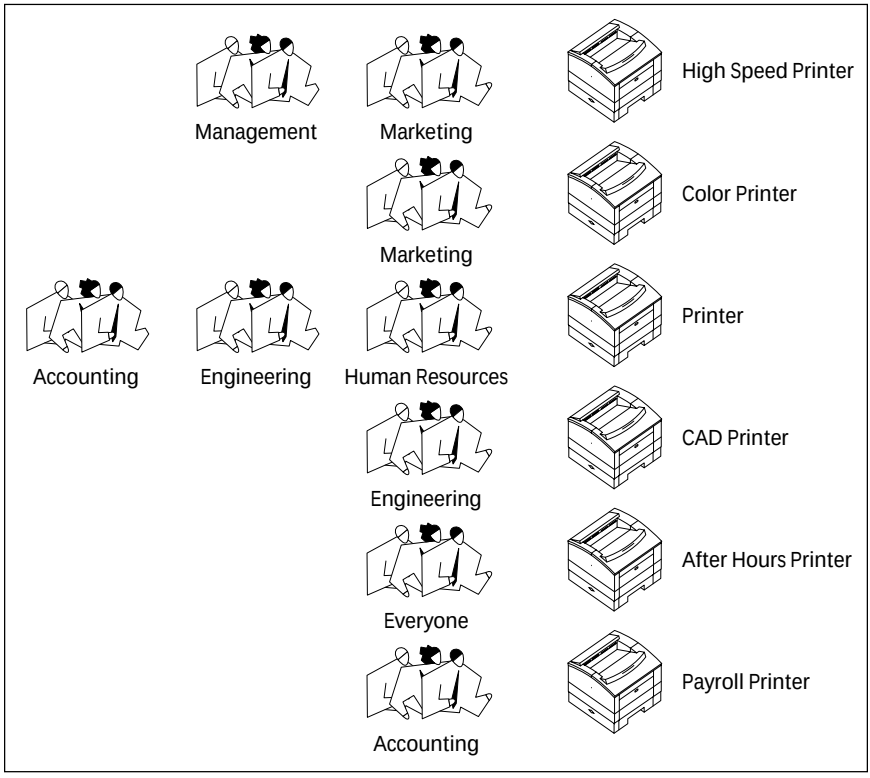

Figure 9.6 Users gain access to resources through their group memberships.

network. For instance, a large number of unsuccessful logon attempts over a very short period of time may indicate that an unauthorized user is attempting to gain access to your network. Other auditing functions provide the basis for informed decisions, such as where to place network resources or what additional resources may be necessary in the future. Windows NT comes with powerful auditing capabilities and allows the auditing of either the success or failure (or both) of the following:

- Logon And Logoff A user logged on or off, or made or broke a network connection to a server or to the local server.
- File And Object Access A user accessed a directory, file, or printer that is set for auditing. This event must be selected to audit file or print resources.

- Use Of User Rights A user exercised a right (does not include logon or logoff).
- User And Group Management A user account or group was created, changed, or deleted; or a user account was renamed, disabled, or enabled, or a password was set or changed.
- Security Policy Changes A change was made to the user rights, audit, or trust relationship policies.
- Restart, Shutdown, And System A user restarted or shut down the computer, or an event has occurred that affects system security or the security log.
- Process Tracking Tracks information for various events, such as program activation.

For the Networking Essentials exam, remember that you must enable auditing before you can track user actions. You must enable File And Object Access as shown in Figure 9.7 if you want to track the use of printers, files, or directories on the server.

Diskless Workstations

Diskless workstations have been in the new lately—more than they have ever been—with the hype surrounding the network computer. However, diskless workstations have been around for quite some time. They can be used in high-security environments, where data compromise is not allowed. With the advent of the Java operating system, diskless workstations (renamed "network computers") are being designed and marketed to appeal to a broader audience. Because these computers only run applications

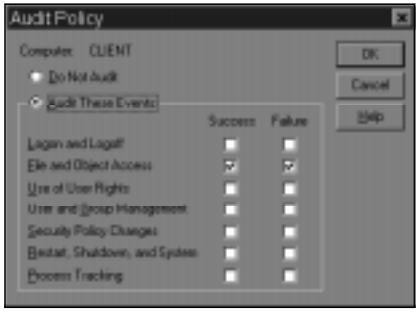

Figure 9.7 Windows NT allows you to enable auditing through the Windows NT User Manager application.

downloaded from the server, they could save maintenance costs associated with updating and installing applications on every computer in an organization. It is still undetermined if network computers will be accepted as viable replacements for the standard personal computer.

A diskless computer has a ROM boot chip that allows the client to initiate a session with the server. At startup, diskless computers broadcast their physical hardware address to the network. The Reverse Address Resolution Protocol (RARP) server does a reverse lookup of the client's assigned network address and sends boot information directly to the client. In addition, the server directs the client to challenge the user. Once the user provides proper authentication, the server allows the client to enter the network.

Diskless computers are sound investments for high-security environments. Because such computers do not have floppy or hard disks, users cannot download data and leave the premises. They are, however, useless if the network server is unavailable.

Encryption

Encryption is the act of changing a message into a format that only the sender and receiver can understand. Encrypting data prior to it entering the network interface card is the most secure means of sending data. Once the data reaches the intended recipient, it is unencrypted using a previously agreed-upon code, and returned to its original format.

The Data Encryption Standard (DES) is an encryption system created for the U.S. government. To use this system, both sender and recipient need access to complementary set of keys. To put it in simple terms, the sender uses a key to encrypt the data before it is sent, and the recipient uses a matching key to unencrypt the data. DES is the de facto standard for data encryption, but is vulnerable if an attacker can obtain a copy of the keys.

Pretty Good Privacy (PGP) is an encryption system much like DES (in both systems, the key is based on a numerical algorithm of specified length). However, PGP uses two linked keys. One key is used for encryption, and the other for decryption. The user safeguards the encryption key, and publishes the decryption key. These keys are generated at the same time, using an algorithm that allows one to encrypt and one to decrypt. The keys cannot be used backwards. This allows the sender to obtain the recipient's public PGP key and encrypt a message to the recipient. As the recipient has the only key that can decrypt the message

(remember, the encryption key will not work as the decryption key) no one else can decrypt the message.

PGP encryption may be a more realistic alternative to encrypting every packet that leaves a computer. It allows the user to choose which data will be encrypted (sensitive email, for instance).

For a more secure environment, you may wish to obtain hardware that encrypts and decrypts data as it enters and leaves your computer.

Virus Protection

Although most viruses are only annoying programs that do not seriously harm a computer, some viruses can cause data loss. When we discuss viruses, we are really talking about four types of computer programs:

- Viruses Computer programs that make copies of themselves and "inhabit" an executable file. (Inhabiting a data file is pointless, because the processor does not execute data.)
- Worm programs These only make copies of themselves; they do not inhabit other files.
- Trojan Horse viruses Much like in the original story, computers enter disguised as a totally different program. Only upon execution do we see what they really are. An example of this is the infamous AOL4FREE.EXE Trojan Horse. Many users believed it to be a downloadable America Online upgrade. Actually, on execution, it changed to the root directory of the host computer and, using the DOS DELTREE command, systematically deleted every file from the user's hard drive. Users could easily stop it by pressing CTRL+C, but many did not know this.
- Macro viruses These are a special subset of the virus family. Using the macro languages associated with many of today's powerful applications, it is possible to create a virus that infects data files. Although the data file itself is never processed, the macro code associated with the data file is. The Concept virus is an example of a macro virus that can spread through a data file attached to an email message.

Virus protection programs cannot prevent viruses. They can only prevent viruses from activating, correct damage done by a virus, and remove viruses from the system. The key to preventing viruses from activating is to have a multiphased plan that addresses the following issues:

- What to do to prevent viruses from entering the network. Limit user-loadable software, and create a strict policy in which all disks are scanned for viruses prior to being used.
- How to check the network for any viruses. Use a virus-checking program regularly to diagnose any inactive viruses that could be on your network. Ensure the program is updated regularly to reflect new viruses.
- What to do when a virus is discovered. Ensure that users know what to do if a virus is discovered. They should inform the Information Management staff, clean their hard drives, clean any other floppies, and inform their peers who may have received an infected disk or email.

Disaster Recovery

Although it is much more exciting to discuss data security and protection in terms of defending your network from intruders, in reality, it is much more likely that you will lose data due to some type of hardware failure. How many times have you thrown away a floppy that just wouldn't format properly or was acting strangely? Or, even worse, it was the file that had the big presentation on it, and it was the only available copy. No, a hacker probably didn't hack your disk. It was a physical object that had a mechanical failure.

Can you imagine what would happen if, despite all the warnings, someone in your office smoked a cigarette in the building, causing the sprinkler system to go off. It may seem amusing at first, until you remember that the sprinkler nozzle you'd been meaning to have the building superintendent remove from the server closet was fully functioning. After getting the key, you already know the network is down. And it will only come back up once you repair the damage caused by the sprinkler.

Data loss can have a disastrous effect on your organization. You are responsible for planning and implementing your "Disaster Policy." It should cover problems caused by fire, data deletion and corruption, theft or vandalism, power fluctuations, server component failure, and natural disasters.

You can respond to most disasters with these three general techniques:

- Tape backup
- Uninterruptible Power Supply (UPS)
- Fault tolerance

Tape Backup

What are the first three rules of computing? Backup, backup, and backup. You may have heard this joke before, but its implications are absolutely clear. An administrator who does not back up mission-critical data is tempting fate. It is the easiest and least expensive means to protecting your data.

Planning a backup strategy is essential in every network. The plan must include what is deemed essential data, which must be backed up more often, and less essential data, which need not be backed up as often. Just as essential as planning and implementing backups is the act of testing the integrity of your system by performing regular test restores

Components of regular network data backups are the hardware, the backup schedule, and the person(s) responsible for performing the backup.

Backup Hardware

Most backup hardware uses backup tapes; however, magneto-optical drives are becoming more widespread in backup operations. Your hardware will reflect your budget, but it should meet the following criteria:

- Size Can it back up the system effectively?
- Speed Can it back up the system in a timely manner?
- Reliability Can it be trusted to work when the system has crashed (also known as, "would you bet your job on it saving the system?")
- Features Does it include error detection and correction capabilities?

Backup Schedule

There are five types of data backup:

- Differential Backs up selected data that has changed since the last backup, without marking data as backed up.
- Incremental Backs up and marks selected data that has changed since the last backup.
- Copy Backs up all selected data without marking data as backed up.
- Daily Backs up only those files that were changed that day, without marking them as backed up.
- Full Backs up and marks all selected data, regardless of if the data has changed since the last backup.

Depending on the network and the data sensitivity, backups could be performed daily, weekly, or even monthly. A good strategy is to make an incremental backup daily, with a full backup over the weekend. This maximizes backup throughput during hours of low network usage. It is also recommended to rotate the backup tapes. Using the same five tapes, labeled Monday through Friday, over the course of a year is probably not a good idea. You should also keep an additional full backup off site, in a secure location. If a fire destroys the building, the data would still be safe.

Backup Operators

Backup operations should be recorded with the date, time, and type of backup. This information should also be written in a network log. Backup operators should be well trained and dedicated to the important task of performing the system backup. For essential security, anyone who has access to your backup tapes should be well known and trusted.

Uninterruptible Power Supply

An Uninterruptible Power Supply (UPS) is a large switched battery that connects between the wall outlet and the computer. If the AC power should fail, the UPS would detect the power loss and immediately begin supplying the computer with power. This occurs so quickly that the computer doesn't crash due to the power outage. Many UPSes have the ability to inform the computer that it is on battery power, and how long the anticipated battery life is. This allows the computer to calculate when to conduct a complete shutdown prior to total power failure. Most systems also will send administrative alerts when the computer changes from AC to UPS power.

In addition, most UPSes act as line conditioners. The voltage at a wall outlet, although rated at 110 to 120 volts, can fluctuate well above and below these levels. These varying levels can damage delicate electrical components. A UPS that also conditions the line ensures that the computer receives voltage within the proper range.

Many operating systems (including Windows NT) do not immediately write data to disk. The system caches the disk writes, and waits for processor idle time. Upon reaching a certain threshold, the data is written to disk. Although this system is the most efficient way for the system to operate, if the power should fail, all data still in the cache will be lost. The UPS provides the necessary time for the administrator, to shut down the system safely (or for the system itself to shut down).

Windows NT can have problems identifying UPSes during startup. When NTDETECT.COM sends a signal to the serial ports to detect attached hardware, it can cause some UPSes to switch off. To correct this, use the No_Serial_Mice switch in the BOOT.INI file to stop NTDETECT.COM from sending the signal.

Fault-Tolerant Systems

Fault tolerance is a term used to describe the ability of a system to recover from failure. Fault-tolerant hard disks are defined by a series of specifications known as RAID (Redundant Array of Inexpensive Disks). The six RAID levels discussed in this section are:

- RAID level 0 Disk striping without parity
- RAID level 1 Disk mirroring or duplexing
- RAID level 2 Disk striping with ECC
- RAID level 3 Disk ECC stored as parity
- RAID level 4 Disk striping with large blocks
- RAID level 5 Disk striping with parity

Each of these levels describe how multiple physical hard disks can be used to increase system performance and/or fault tolerance, in addition, we include information on a concept called sector sparing.

RAID Level 0—Disk Striping Without Parity

Strangely enough, stripe sets (without parity) provide no fault tolerance. If any disk fails, all data is lost. However, disk throughput is high on a system that uses disk striping without parity. This occurs because the slowest element in the sequence of writing to disk is the hard drive. With two drives in a stripe set, one drive will receive data from the controller and begin writing to disk. The controller can immediately switch to the next disk and send it data, and that disk can begin saving the data. The first disk may not have completed the first write operation, and the controller may have to wait. As you increase the number of hard drives in the stripe set, more time is available for each drive to write to disk.

RAID level 0 provides NO data redundancy. Its attraction lies in its ability to conduct read and write operations more efficiently than the other RAID levels.

RAID Level 1—Disk Mirroring Or Duplexing

RAID level 1 allows an operating system to save data to two separate hard drives. This level is normally associated with disk mirroring and disk duplexing. Disk mirroring uses two hard drives and one hard-drive controller. If one drive fails, the other hard drive is available with a complete set of all the data. Disk duplexing employs two hard drives, each of which has its own hard-drive controller (see Figure 9.8). This technique also provides redundancy for the hard-drive controllers. As in disk mirroring, if one drive fails, the other would still be available. However, in disk duplexing, if a hard-drive controller fails, or if a disk fails, all data is still available. A drawback to RAID level 1 is low-disk utilization. With two disks in a set, disk utilization is only 50 percent. If there are two 1-GB drives in a system with RAID level 1 active, there will only be 1 GB available for use.

Windows NT Server supports both disk mirroring and disk duplexing.

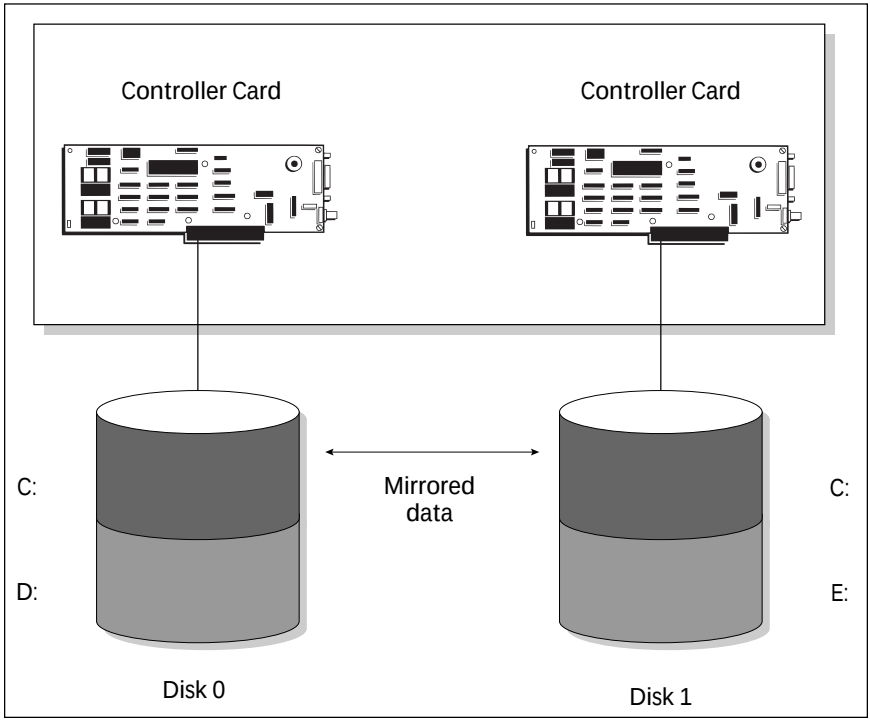

Figure 9.8 Disk duplexing is similar to disk mirroring, except each mirrored disk has its own controller.

Disk mirroring uses only one controller for two hard drives, while disk duplexing uses a controller for each hard drive. Also, disk utilization for RAID level 1 is 50 percent.

RAID Level 2—Disk Striping With Error Correction Code

RAID level 2 implementations write data across a minimum of three disks linked to one controller, and Error Correction Code (ECC) is used to keep track of the data. The group of disks is called a "stripe set," and the process is referred to as data striping. As the data is being written to the stripe set, the ECC code is written to a separate disk. Because ECC uses more space than parity code and only provides marginal improvement in disk utilization, this level is seldom used.

RAID Level 3—Disk Striping With ECC Stored As Parity

RAID level 3 is similar to RAID level 2, except that the ECC is replaced with a parity-checking scheme. Instead of calculating ECC, the system checks the data by adding a mathematical formula to the data, where the number of 1s in the data must always be the same. This data-checking system is known as parity checking. It allows the system to check to ensure the data is correct and uses less disk space than ECC. Only one disk is used to store parity data; if this disk fails, the data may not be recoverable, depending on the last backup.

RAID Level 4—Disk Striping With Large Blocks

RAID level 4 defines the method of writing entire blocks of data on each disk, rather than striping blocks of data across the disks in the set. RAID level 4 striping uses parity information to ensure the data is written correctly. The parity information is stored on one disk. As a result, each time data is written to the stripe set, the parity must be calculated and written to the parity disk. This system works well with the large data blocks, but is not efficient because the parity information must be written as the data is written. Windows NT does not support RAID levels 2, 3, and 4, but they can be supported as a hardware-based RAID implementation.

RAID Level 5—Disk Striping With Parity

This RAID level writes data and parity information across all disks in the stripe set, ensuring that the parity information associated with that data is not on the same disk as the data (see Figure 9.9). If a single disk fails, the parity information (in conjunction with the data still on the other disks)

can replace the missing data dynamically. However, if two disks fail at the same time, the stripe with parity cannot recover the information. The only solution in this instance is to restore from tape. RAID level 5 is supported by Windows NT and is a more efficient technique than levels 2, 3, and 4.

Because parity information must be written with each write operation, some of the drive space is occupied by parity information. This means that the usable drive space is equal to the size of all the drives in the system minus 1 divided by the number of physical disks in the system times the total space of all the drives in the system. For example: If a system is using four 3 GB drives, the total space on the set would be: 3 GB + 2 GB + 1 GB + 6 GB = 12 GB. The parity information would use 25 percent of the stripe set (1 divided by the total number of physical disks [4]), which is 3 GB. The total space available for data in this scenario would be 9 GB (or 12 GB minus 3 GB).

Another example: If a system is using six disks with 2 GB each, the total space on the set would be: 2 GB + 2 GB + 2 GB + 2 GB + 2 GB + 2 GB = 12 GB (same as the earlier example). The parity information would use one sixth of the stripe set (1 divided by 6), which is 2 GB. The total space available for data would be 10 GB (or 12 GB minus 2 GB).

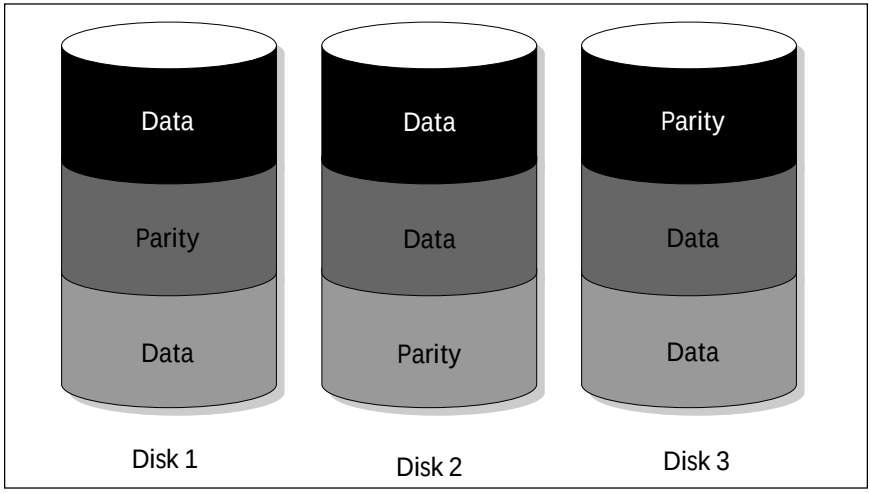

Figure 9.9 Disk striping with parity provides fault tolerance by writing parity information during each write operation—the disk controller writes the parity information in alternating locations in the set, which keeps parity information spread evenly across all disks in the stripe set.

Because the disks in a stripe set with parity must be the same size, the easy way to calculate the total space available for data is to add up the total available space on all of the drives and subtract the total available space on a single drive. As the number of disks in a stripe set increases, more space is available for data storage, because less space is necessary for the parity information.

Sector Sparing

Sector sparing is a fault tolerance technique where the hardware or the operating system checks the integrity of the disks prior to write operations. If it discovers a bad sector during a write operation, the bad sector is mapped out (marked unusable) and the data is written to a good sector. This technique only works on SCSI drives, unless the file system supports sector sparing (as does the Windows NT File System).

Table 9.1 Fault tolerance supported by Windows NT—hardware requirements.

RAID Level	Minimum Number Of Disk Drives	Maximum Number Of Disk Drives	Number Of Controllers
RAID 1: disk mirroring	2	2	1
RAID 1: disk duplexing	2	2	2
RAID 5	3	32	1 or more
RAID 0	2	32	1 or more

Exam Prep Questions

Question 1

The easiest and least expensive means to protect data is:

- ❍ a. Fault tolerance
- ❍ b. RAID level 5
- ❍ c. Disk mirroring
- ❍ d. Scheduled backups

The answer for this question is d. Backups are the easiest and least expensive way to protect data. Although various RAID strategies are very effective at providing fault tolerance, no other means is as easy, or as comprehensive.

Question 2

Which of the following are normally associated with user-level security? [Check all correct answers]

- ❑ a. Each user is assigned a unique user name and password
- ❑ b. Windows 95
- ❑ c. Windows NT
- ❑ d. Users are allowed access to resources based on their user names and passwords

The answers for this question are a, c, and d. Systems with user-level security assign each user a unique user name and password, and assign resources based on privileges associated with the user. Contrast this with share-level security, where resource owners place passwords on resources. Windows NT is designed to support user-level security. Windows 95 is designed to support share-level security.

Question 3

You are consulting for a company that is installing NetBIOS as the only protocol. The network topology consists of a small 10BaseT LAN with 1 Windows NT PDC, 1 BDC, and 15 NT workstations. The company has functional areas named Management, Accounting, Engineering, and Information Systems. The administrator wants the computer name of each workstation to include the user's functional area, the word "workstation" or "server," and the number of the system.

Required result:

Ensure the computers can communicate.

Optional desired results:

The system should be able to resolve FQDNs to IP addresses.

The system should be able to resolve computer names to IP addresses.

Proposed solution:

Use the following system to name the computers:

Accountingworkstation1	Accountingworkstation2
Accountingworkstation3	Accountingworkstation4
Accountingworkstation5	Managementworkstation1
Managementworkstation2	Managementworkstation3
Managementworkstation4	Managementworkstation5
Managementworkstation6	Engineeringworkstation1
Engineeringworkstation2	Engineeringworkstation3
Engineeringworkstation4	Informationserver1
Informationserver2	

Which results does the proposed solution produce?

- ❍ a. The proposed solution produces the required result and produces both of the optional desired results.
- ❍ b. The proposed solution produces the required result and produces only one of the optional desired results.
- ❍ c. The proposed solution produces the required result but does not produce any of the optional desired results.
- ❍ d. The proposed solution does not produce the required result.

The correct answer is d. NetBIOS only allows 15 characters as computer names, with a 16th character used by the system. The system could easily be implemented by removing the word "workstation" or "server" from the proposed computer names. DNS resolves FQDNs to IP addresses. WINS resolves computer names to IP addresses.

Question 4

Disk striping without parity requires how many disks to implement?

- ❍ a. One
- ❍ b. Two
- ❍ c. Three
- ❍ d. RAID 0

The correct answer is b. Disk striping without parity only requires two disks to implement. Disk striping with parity requires a minimum of three disks. Both techniques can use up to 32 disks in Windows NT.

Question 5

Fault tolerant systems include which of the following? [Check all that apply]

- ❑ a. Disk duplexing
- ❑ b. Sector sparing
- ❑ c. Data encryption
- ❑ d. RAID level 0

The correct answers are a and b. Disk duplexing is a fault-tolerant method that uses two disk controllers and two disk drives. Although sector sparing is not a RAID level, it does supply fault tolerance to systems by detecting bad sectors, attempting to save the data in the bad sectors to a good sector, and identifying the bad sector as unusable. Data encryption secures data from being read; it does not secure it from being lost or damaged. RAID level 0 provides no fault tolerance.

Question 6

Your UPS shuts off every time your Windows NT computer boots. What should fix the problem and still maintain protection from power failure? [Choose the best answer]

- ❍ a. Bypass the UPS and plug your computer directly into the wall outlet.
- ❍ b. Push the spacebar to invoke The Last Known Good Configuration.
- ❍ c. Use the No_Serial_Mice switch in the BOOT.INI file.
- ❍ d. Reboot using the Emergency Repair Disk.

Answer a does nothing to protect your machine from a power failure; therefore, answer a is incorrect. If the problem is recurrent, ten the LKGC won't fix the problem; therefore, answer b is incorrect. The correct answer is c. During boot, NTDETECT.COM attempts to identify all system peripherals. This includes sending a signal to the serial ports. This signal may shut off some manufacturers' UPSs. Like answer b, the ERD won't fix the problem; therefore, answer d is also incorrect.

Question 7

A company is installing an Internet Web server and is very concerned with how quickly data images can be read from the server and sent to potential clients. What disk management system would you want to enable? The company is not as concerned with losing the data as it is about the speed at which the data can be read.

- ❍ a. One
- ❍ b. RAID level 2—disk striping with ECC
- ❍ c. Disk mirroring
- ❍ d. Raid level 0—Disk striping without parity

The correct answer is d. Disk striping without parity is much quicker than the other choices at Read and Write operations because no parity information is calculated or written to disk.

Question 8

A user wants to share a directory on his computer with coworkers, and does so using a password to protect the contents of the directory. This is an example of:

- ❍ a. User-level security
- ❍ b. Share-level security

The correct answer is b. Users assigning passwords to resources is an example of share-level security. Access to resources depending on rights that the administrator assigns is an example of user-level security.

Question 9

You are consulting for a company that is using TCP/IP as the only protocol. The network topology consists of a small 10BaseT LAN with 1 Windows NT PDC, 1 BDC, and 15 NT workstations. Due to the nature of some of the data, the company wants to install some type of software-based fault-tolerant system.

Required result:

Ensure the data is immediately available in case of single disk failure.

Optional desired result:

Ensure the system is inexpensive.

Proposed solutions:

Install disk striping without parity.

Schedule and conduct frequent data backups.

Which results do the proposed solutions produce?

- ❍ a. The proposed solutions produce the required result and produces the optional desired result.
- ❍ b. The proposed solutions produce the required result but do not produce the optional desired result.
- ❍ c. The proposed solutions do not produce the required result.

The correct answer is d. Although the backups are present, the data would not be immediately available in the case of a single disk failure. The administrator would have to restore the backup data before it could be used. This would be an excellent solution if the company installed disk striping with parity. RAID level 5, disk striping with parity, is software that Windows NT supports. A hardware-based RAID solution would not need to be purchased.

Need To Know More?

Chellis, James, Charles Perkins, and Matthew Strebe: *MCSE: Networking Essentials Study Guide*. Sybex Network Press, San Francisco, CA, 1994. ISBN 0-7821-1971-9. Chapter 8, "Administering Your Network," contains excellent coverage of the various topics contained within this chapter.

Microsoft Press: *Networking Essentials*. Redmond, Washington, 1996. ISBN 1-55615-806-8. Unit 6, Lesson 23, "Avoiding Data Loss," discusses all of the fault-tolerance topics in this chapter in greater detail.

Rutstein, Charles B.: *National Computer Security Association Guide to Windows NT Security: A Practical Guide to Securing Windows NT Servers & Workstations*. McGraw-Hill, 1997. ISBN 0-07-057833-8. This is the best guide to NT security on the market.

Search the TechNet CD (or its online version through www.microsoft.com) using the keywords "RAID," "fault tolerance," "WINS," and "DNS."