

9

Introduction to the Server

After studying this chapter, you will be able to:

- Recall the various roles of servers in a network environment.
- Give examples of uses for a thin server.
- Give examples of uses for a thin client server.
- Differentiate between a server and a PC.
- Identify the ID and LUN numbers of a Small Computer Systems Interface (SCSI) device.
- Identify the four types of system resources.
- Compare RAID systems.
- Recall the characteristics and purpose of network attached storage (NAS) technology.
- Recall the characteristics and purpose of storage area network (SAN) technology.
- Give examples of common Fibre Channel technology applications.
- Identify the Fibre Channel topologies.
- Interpret an iSCSI IQN address.

Network+ Exam—Key Points

The Network+ Certification exam specifically references fault tolerance and network storage. Pay particular attention to the characteristics of RAID 0, RAID 1, RAID 5, and RAID 6 and to network storage, such as storage area network (SAN) and network attached storage (NAS).



Key Words and Terms

The following words and terms will become important pieces of your networking vocabulary. Be sure you can define them.

ACPI Component Architecture (ACPICA)	Internet Small Computer Systems Interface (iSCSI)
Automatic Configuration and Power Interface (ACPI)	interrupt request (IRQ)
backplane	logical unit number (LUN)
blade server	memory address assignment
Direct Memory Access (DMA) channel	Network Address Authority (NAA)
disk mirroring	network attached storage (NAS)
disk striping	parallel processing
duplexing	power-on self-test (POST)
error correction	Redundant Array of Independent Disks (RAID)
fabric switch	serial attached SCSI (SAS)
fault tolerance	Small Computer Systems Interface (SCSI)
Fibre Channel	storage area network (SAN)
firmware	thin client
hot swapping	thin client server
hot-swap technology	thin server
Input/Output (I/O) port	

A complete understanding of server hardware can take a great deal of study. This chapter provides an overview of the server and identifies different server types and their roles. You will be introduced to some of the hardware that makes the server unique from the ordinary PC, and you will learn about RAID and other storage systems.

Server Types and Services

Servers provide a variety of services. Some of the services a server can provide are authentication and security, Web, mail, and print. A server can be called by many names. For example, it can be called an *authentication and security database server*, *Web server*, *mail server*, and *print server*, **Figure 9-1**. A network may have a single server that provides a variety of services, or it may have a group of servers, each providing a specific service.

A small network usually has one server set up to handle many different services. A large network usually has several servers, each providing a different service or set of services. For example, a large corporation may use one server to handle e-mail requests and Web hosting, another server to serve as a domain




Figure 9-1. Servers provide a variety of services, such as authentication and security, Web, mail, and print.

controller to provide security for the entire network, and another server to provide application software, a database for its clients, and support for print operations.

The more services a server provides and the more clients it services creates a higher demand on the server. A single server with a limited amount of services may be fine for a relatively small number of users, typically less than 25. A commercial enterprise spanning across the country or world may require hundreds of servers. The network administrator or network designer decides the number and capacity of individual servers needed by taking into account the number of users and the system's predicted network traffic.

Each network system is uniquely designed, even though each network has many similarities. Some network equipment providers have software programs that help you design a network. You simply enter information, such as the number of clients, offices, cities, and countries and the type of software and services to be provided. After all the information is collected, the software program provides an estimate of the size and number of servers required. We will discuss this topic later in Chapter 19—Designing and Installing a New Network.

Thin Servers

A *thin server* is a server that has only the hardware and software needed to support and run a specific function, such as Web services, print services, and file services. It is more economical to use a thin server as a print server than to tie up a more expensive server simply to handle printing on a network. IBM markets a thin server, which consists of a sealed box that contains only the essential hardware and software required for supporting the server's dedicated function.

Thin Client Servers

A *thin client server* is a server that provides applications and processing power to a thin client. Thin client servers run terminal server software and may have more RAM and hard disk drive storage than needed. A *thin client* is a computer that typically has a minimal amount of processing power and memory.

thin server
a server that has only the hardware and software needed to support and run a specific function, such as Web services, print services, and file services.

thin client server
a server that provides applications and processing power to a thin client.

thin client
a computer that relies on the thin client server's processing power and memory.

It relies on the server's processor and memory for processing data and running software applications. Thin clients are becoming common in industry for such applications as hotel bookings, airline ticketing, and medical record access.

When budgets are tight, you can install what is normally considered an obsolete PC on a thin client network. The obsolete PC is obsolete because of its processing speed, lack of storage, and inability to run new software. The thin client server can provide all the services that are required by the obsolete PC, thus making the obsolete PC a useful workstation.

Do not confuse a thin client with a dumb terminal. A dumb terminal sends user input to a mainframe. Dumb terminals have absolutely no computing power, operating system, hard disk drive, BIOS, and RAM. A thin client may not need a hard disk drive or an operating system; however, it is still a full-fledged computer because it has a CPU and processing power.

Windows 2000 Server, Windows Server 2003, and Windows Server 2008 come with Terminal Services software. Once Terminal Services is set up on the server, you have the option to generate a disk that will assist you in automatically setting up thin client workstations.

Server Classification by Number of CPUs

Servers vary a great deal by size and power. The number of processors they contain usually classifies them. For example, Sun Microsystems has three classifications of servers: entry-level, midrange, and high-end. Sun Microsystems defines these classifications by the number of processors they contain. A Sun Microsystems entry-level server has up to 8 processors. A midrange server has up to 30 processors, and a high-end server has up to 106 processors. Other vendor's definitions vary somewhat, but this gives you a general idea. In this chapter, we focus on the HP ProLiant DL740 server. The HP ProLiant DL740 is one of the most powerful entry-level servers manufactured.

Major Server Components

It is assumed you have some PC hardware background from completing a course on PC service and repair or have a CompTIA A+ Certification. If not, it is strongly advised that you take a course in PC repair or, at the very least, complete a home study program. A good background in PC hardware proves to be beneficial to anyone working with network servers. As a matter of fact, IBM and other companies require the CompTIA A+ Certification as part of its progression toward certification as a server technician.

When handling electronic components, it is standard practice to use electrostatic discharge (ESD) safety practices. Electronic components based on complimentary metal oxide semiconductor (CMOS) technology can be easily damaged by static electricity. The human body and clothing can build up a tremendous static electricity charge. When touching a static-sensitive electronic component, the static electricity can discharge and cause the component to be destroyed. To avoid damaging electronic components, use an antistatic wrist strap, which is designed to drain static charges safely.

Caution



Do *not* use an antistatic wrist strap when servicing any component with live voltage present.

Do *not* open a CRT (monitor) while wearing an antistatic wrist strap, even if the CRT is unplugged. CRT screens can hold a high voltage charge for long periods of time after they have been disconnected from electrical power.

Caution



The major components of a server are similar to a typical PC. In fact, you can use a typical PC as a server for a small network. A network server can be as simple as a typical PC or as complex as a piece of equipment designed exclusively for networks. Some server models with multiple CPUs, large amounts of RAM, and vast amounts of storage space can be thought of as a "small" mainframe rather than as a PC. There are many server designs available. As part of your study of network servers, it is highly advisable to check the IBM, Sun Microsystems, HP, and Dell Web sites to see the available designs. This section looks at the following components: case, hot-swap components, power supply, motherboard, BIOS, and CPU.

Case

Server case styles vary greatly. Some are similar to desktop models. Some are designed to be mounted into a rack like the server in **Figure 9-2**.

Large enterprise servers are actually a group of individual servers mounted in a rack system or cabinet, **Figure 9-3**. The group of servers acts as one unit that has multiple CPUs and a large amount of system resources, such as hard disk drive storage and RAM.

A *blade server* derives its name from its size and shape. It is extremely thin compared to other servers and fits into a rack, **Figure 9-4**. Blade servers are especially designed to allow a large number of servers to be mounted in a small space. Blade servers are not to be confused with thin servers. A blade server can contain more than one processor and be quite powerful. They are often selected for applications where there is a requirement for many servers to perform the same or similar function. For example, they may be used in a large Web service facility or a file server farm.

blade server
a powerful server that is extremely thin. It is designed to be mounted in a small space with other blade servers.

Hot-Swap Components

Servers are generally designed to provide continuous service with minimal interruptions and data loss. The main feature that most servers use to provide this continuous service is hot-swap technology. *Hot-swap technology* allows a component to be removed or installed while the system is running. There is no need to power down the system while replacing or adding major components. Some common hot-swap components are hard disk drives, memory modules (RAM), CPUs, and power supplies.

Never assume a component is hot-swappable. Always check the system manual before removing any component while the system is powered on.

hot-swap technology
a technology that allows a component to be removed or installed while the system is running.



Figure 9-2. The ProLiant DL740 can be mounted into a rack. (Hewlett-Packard Company)

Figure 9-3.

This cabinet can store many servers, allowing them to act as a single unit. (Hewlett-Packard Company)

**Figure 9-4.**

A blade server is thin, yet powerful. It is designed to allow many blade servers to be installed in a small area. A—A single blade server. B—Many blade servers installed in a single rack. (Hewlett-Packard Company)

**Figure 9-5.**

Back view of the ProLiant ML350. This system comes with two power supplies that evenly distribute power throughout the system. When one power supply fails, the other carries the full load. (Hewlett-Packard Company)



Power Supply

The power supply converts standard 120-volt or 240-volt AC power into lower DC voltage levels that can be used by the motherboard and other devices inside the computer case. A typical power supply provides an assortment of output plugs. The plug style matches the intended piece of equipment. For example, a standard ATA drive uses a Molex plug, which supplies 5 volts DC and 12 volts DC to the ATA drive. An ATX-style motherboard plug supplies 3.3 volts, 5 volts, and 12 volts as well as a ground connection to the motherboard. It is also designed to plug into the motherboard correctly. Older style connections could be incorrectly plugged into the motherboard if you did not pay attention to the color-coding.

Entry-level servers are typically equipped with two power supplies, **Figure 9-5**. The two power supplies balance the load. If one power supply fails, the other power supply carries the full load. Some servers are designed with “hot-swappable” power supplies, which permit the defective power supply to be changed while the server is running.

To prevent the loss of data, most server manufacturers recommend disconnecting the server from the network while replacing hot-swappable power supplies.

Motherboard

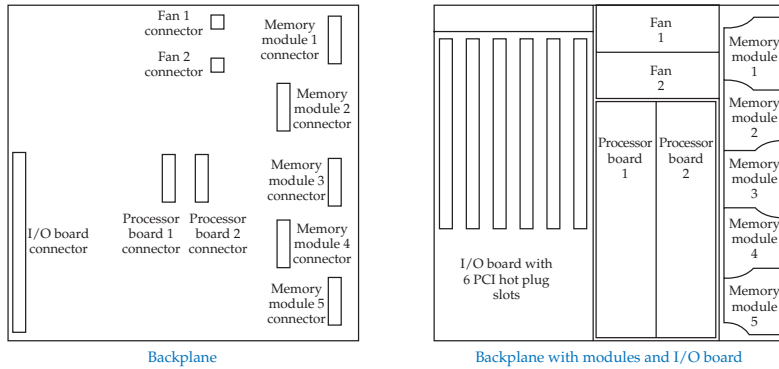
Servers generally use a backplane. A *backplane* is a simple motherboard designed with minimal components. It typically serves as the interface of all the major components. It is designed to allow major components to be added or removed without powering down the system. Removing components without shutting down the system is called *hot swapping*.

Figure 9-6 shows the general backplane layout of the HP ProLiant DL740 server with and without the added modules. An I/O board attaches to the backplane. It contains six PCI hot-swappable slots. The processor boards each contain 4 CPUs, for a total of 8 CPUs. Each memory module in the system is hot-swappable and contains 8 DIMMs. The HP ProLiant DL740 can contain up to 40 GB of RAM. The multiple CPUs and vast amount of RAM are required to provide the many different services to

backplane a simple motherboard designed with minimal components. It typically serves as the interface of all the major components.

hot swapping the process of removing components without shutting down the system.

Figure 9-6.
The backplane layout of the ProLiant DL740.



numerous network clients with minimal delay. The CPUs need not be the latest, high-speed processors available because there are numerous processors in the server.

BIOS

The BIOS contains a small software program that starts the server boot operation when power is applied to the server. The combination BIOS chip and the software program are usually referred to as *firmware*. The BIOS is responsible for the *power-on self-test (POST)*, which is performed at startup. The POST does a quick, initial check of the major components, such as memory, disk drives, keyboard, mouse, and monitor, to be sure that a minimum working system is available. After a general check is made of the major hardware components, the BIOS turns control over to the operating system software. The operating system completes the boot process by loading more advanced hardware drivers than the BIOS did. It also performs a more sophisticated check of hardware and software systems. If all the software and hardware components appear to be in working order, the final screen, which serves as a user interface, appears.

Central Processing Unit (CPU)

One major difference between a server and a typical PC is that many servers have multiple CPUs installed rather than just one. When multiple CPUs are installed, the server can perform *parallel processing*. This means that a program can be processed through more than one CPU simultaneously. Another advantage of multiple CPUs is several clients may be serviced at the same time rather than waiting their turn to access a single CPU. Supercomputers and enterprise servers may contain over 1,000 CPUs.

A small office server can get by with one processor, especially if demand on the server is low. A large enterprise system requires multiple processors to meet the demands of the server's clients. Some processor modules are hot-swappable. When a single CPU fails, it can easily be removed and replaced while the server is running and providing services to clients.

firmware
the combination of BIOS chip and the software program within the chip.

power-on self-test (POST)
a BIOS routine that performs a series of hardware checks to determine if the computer is in minimal working order.

parallel processing
processing a program through more than one CPU simultaneously.

Small Computer Systems Interface (SCSI)

Small Computer Systems Interface (SCSI), pronounced *skuzzy*, is a computer bus technology that allows you to connect multiple devices to a single controller. The benefits of SCSI technology can be readily observed in the chart in **Figure 9-7**. SCSI technology not only allows multiple devices to connect to a single controller, it also supports high-data transfer rates. This is quite a performance improvement when compared to the traditional IDE or ATA attachment. SCSI is ideal for servers that must hold large amounts of data that is accessed by numerous clients.

Most SCSI drives are hot-swap devices. Often, SCSI drives are arranged in the server with hot-swap bays open to the outside of the case to provide easy access. This allows you to connect or disconnect a drive without opening the server case. See **Figure 9-8**.

SCSI technology has evolved over the years. With this evolution, new names have emerged to describe the improved technology. The term *Wide* is used to indicate 16-bit data transfers in place of 8-bit data transfers. To reflect the increase in frequency, the term *Fast* was used and then the term *Ultra*. Combination of the words, such as *Ultra/Wide*, are also used to express the newer technologies. Ultra/Wide means the SCSI device is faster and supports 16-bit transfers. A close

Small Computer Systems Interface (SCSI)
a computer bus technology that allows you to connect multiple devices to a single controller.

Common Name	Class	Devices	Bus Width in bits	Speed	MBps
SCSI-1	SCSI-1	8	8	5 MHz	4-5 MBps
Wide SCSI	SCSI-2	16	8	5 MHz	10 MBps
Fast SCSI	SCSI-2	8	8	10 MHz	10 MBps
Fast/Wide SCSI	SCSI-2	16	16	10 MHz	20 MBps
Ultra SCSI	SCSI-3	8	8	20 MHz	20 MBps
Ultra/Wide	SCSI-3	8	16	20 MHz	20 MBps
Ultra2	SCSI-3	8	8	40 MHz	40 MBps
Ultra2/Wide SCSI	SCSI-3	16	16	40 MHz	80 MBps
Ultra3 SCSI	SCSI-3	16	16	40 MHz	160 MBps
Ultra Wide 320	SCSI-3	16	16	80 MHz	320 MBps
Ultra Wide 640	SCSI-3	16	16	160 MHz	640 MBps

Figure 9-7.
SCSI technology specifications.



Figure 9-8.
Hot-swap drives are typically accessible from the front of the server for easy access. (Hewlett-Packard Company)

study of the table in Figure 9-7 will help you understand the evolution of the SCSI technology. To learn more about SCSI technology, visit the Adaptec Web page at www.adaptec.com. There you will find in-depth knowledge about SCSI systems and other technologies.

A SCSI hardware system consists of a host adapter or controller card, a flat ribbon cable to connect SCSI devices, and SCSI devices such as disk drives, CD-ROMs, and tape drives. SCSI cables come in a variety to match the many different classifications of SCSI. The cables come as DB-25, 50-pin, 68-pin, and 80-pin styles.


Look at **Figure 9-9** to see the way a typical SCSI host adapter, cable, and devices might appear. The host adapter and devices can connect to any part of the chain. The chain of SCSI devices must be terminated. Termination is typically completed at the last device with a termination block.

SCSI ID numbers

Each device, including the host adapter card, must have a unique ID number. The numbers start at zero and end at seven for an eight-device SCSI chain and end at fifteen for a sixteen-device SCSI chain. The host adapter typically is assigned the highest number.

SCSI ID numbers are assigned through pins and jumpers on the SCSI drive. The pins are arranged in pairs. When a jumper is applied across the pins, an

Figure 9-9. Typical SCSI drive arrangement.



electrical connection is made. The pairs of pins represent the binary number system. The jumpers are applied in a binary pattern that represents the SCSI ID number. **Figure 9-10** illustrates each of the binary patterns and the related SCSI ID number.

Logical unit number (LUN)

SCSI devices are not limited to internal devices. The SCSI chain can extend outside the case by adding an extender card to the SCSI chain. The extender card allows more devices to be attached to existing SCSI systems. The extender card is an integrated circuit card connected to the SCSI chain as a SCSI device. See **Figure 9-11**.

When additional devices are attached, they are identified separately from the original chain of devices. Each additional device connected to the SCSI extender is identified with a **logical unit number (LUN)**. SCSI bus extenders are also referred to as *SCSI expanders, repeaters, and regenerators*.

External SCSI device chains are commonly used in RAID systems. RAID systems are covered later in this chapter.

Logical unit number (LUN) a numbering scheme to identify SCSI devices attached to an extender card.

Serial Attached SCSI

Serial attached SCSI (SAS) is the next generation of SCSI. It uses a serial port similar to the SATA serial drive port, **Figure 9-12**. This eliminates the need of the wide SCSI cable and produces faster data transfer rates. First generation SAS devices can achieve data rates as high as 3 Gbps. Second generation SAS devices can achieve rates of 6 Gbps. Note that these speeds are expressed in bits per second (bps), not bytes per second (Bps). This is because SAS transfers data in a serial fashion. The SAS design allows 128 devices to be attached directly.

Serial attached SCSI (SAS) SCSI technology that transfers data in serial fashion.

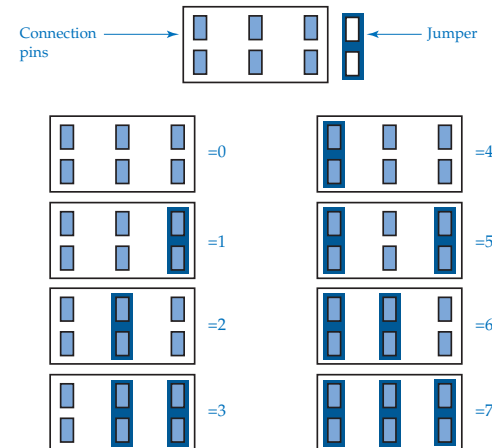


Figure 9-10. SCSI binary patterns.

Figure 9-11. The SCSI limit of seven devices can be expanded by using expander cards. Each expander card allows an additional seven devices to be connected to the SCSI system. LUN numbers are assigned to the additional SCSI devices so they can be identified by the system software and BIOS.

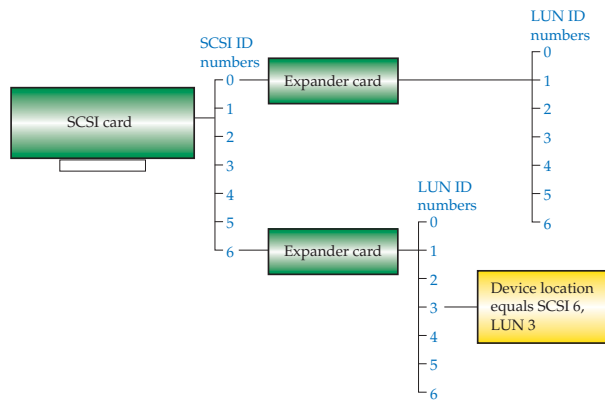



Figure 9-12. SAS and SATA disk drive connectors. Notice the two connector types are similar.



Through the use of edge expander devices, a maximum of 16,000 SAS devices can be connected.

ACPI and ACPIA

Automatic Configuration and Power Interface (ACPI) a standard for controlling the power options for computer hardware devices.

Automatic Configuration and Power Interface (ACPI) is a standard for controlling the power options for computer hardware devices. The first implementation of the ACPI standard gave users the ability to select how they wanted their computer to behave to conserve power. For example, the user could configure the computer or individual computer devices to go into a state of hibernation after a period of inactivity.

Figure 9-13 shows the power options available for a wireless network adapter. These options allow the computer to turn off the adapter to save power and allow the adapter to wake the computer when network activity occurs.



Figure 9-13. Power options for a wireless network adapter.

ACPI Component Architecture (ACPICA) is an extension to the original ACPI standard. The goal of ACPICA is to create a non-proprietary software package for configuring hardware. An ACPICA-compliant device is operating system independent. ACPICA is written in the C language and is an open standard allowing any computer software or hardware manufacturer to share power-saving features rather than developing proprietary power-saving features. ACPICA is compatible with both 32-bit and 64-bit systems.

Some of the companies that participate with the goals of ACPICA are Microsoft, HP, Intel, and several Linux organizations. To learn more about ACPICA, check out the ACPICA Web site (www.acpica.org).

ACPI Component Architecture (ACPICA) an extension to the original ACPI standard that provides a non-proprietary software package for hardware configuration.

System Resources

System resources refer to resources such as interrupt requests, Direct Memory Access channels, Input/Output ports, and memory. System resources are assigned to components installed in the computer system such as hard disk drives, keyboards, and mice. In a Windows-based server, you can view system resource assignments in Device Manager.

Not all devices have all four types of resources assigned to them. For example, some devices do not require a DMA channel assignment. System resources typically cannot be shared between two devices. Sharing the same system resource causes a system resource conflict, referred to simply as a *conflict*. Conflicts must be resolved before the system can operate properly. An exception to this rule is motherboard chips used to bridge different bus systems. A motherboard chip can share an IRQ with a hardware device.

Interrupt Request (IRQ)

An **interrupt request (IRQ)** is a circuit that communicates with the CPU. Hardware devices send an electrical signal to the CPU using an assigned IRQ circuit. There are 16 IRQ assignments numbered from 0 to 15, **Figure 9-14**. Many of the assignments cannot be changed, but some can. When Plug and Play devices are used, the IRQ is automatically assigned. Typically, each hardware device must use a separate

interrupt request (IRQ) a circuit that communicates with the CPU.

Figure 9-14.
IRQ assignments.



IRQ to communicate with the CPU. If two hardware devices are assigned the same IRQ, an IRQ conflict occurs. Once the first device contacts the CPU using the IRQ assignment, the other device cannot communicate with the CPU. An IRQ conflict can lead to problems such as an inoperable device, a system crash, and a system lockup. An IRQ can typically be assigned manually by first entering the BIOS setup program and disabling Plug and Play detection. Changes to the IRQ can then be made through Device Manager. After the problem device has been assigned the proper IRQ, the Plug and Play detection can be activated without a problem.

You may notice IRQ numbers higher than 16 in Device Manager. These numbers are virtual IRQ settings, rather than hardware-based IRQs. They work the same as the hardware-based interrupts, but they are not limited by physical properties. Virtual IRQ assignments are configured automatically by the operating system.

Tech Tip



Yellow question marks next to devices in Device Manager indicate there is a problem with the device or assignment.

Direct Memory Access (DMA) Channel

Direct Memory Access (DMA) channel
a circuit that allows devices to communicate and transfer data to and from RAM without the need of CPU intervention.

A *Direct Memory Access (DMA) channel* is a circuit that allows devices to communicate and transfer data to and from RAM without the need of CPU intervention. Computer systems were first designed with the CPU handling all communication. Every bit of data had to travel through the CPU. DMA technology was introduced to save valuable processor time.

DMA works in conjunction with the chipsets on the motherboard, BIOS software, and the CPU. Large blocks of data that need to be transferred between hardware devices and memory are transferred through a DMA channel that is assigned to the device. **Figure 9-15** shows the DMA controller identified by Device Manager.

One of the devices controlled by the DMA memory access controller is the computer ATA drive. **Figure 9-16** shows the **Advanced Settings** tab for the IDE channel. The IDE channel has been assigned to a DMA channel to take advantage of fast data transfer.

Figure 9-15.
DMA assignment.

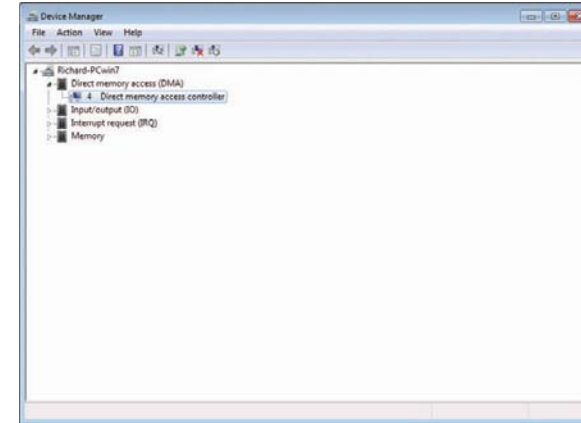



Figure 9-16.
IDE Channel Properties dialog box **Advanced Settings** tab. The IDE channel is assigned to a DMA channel by enabling DMA.



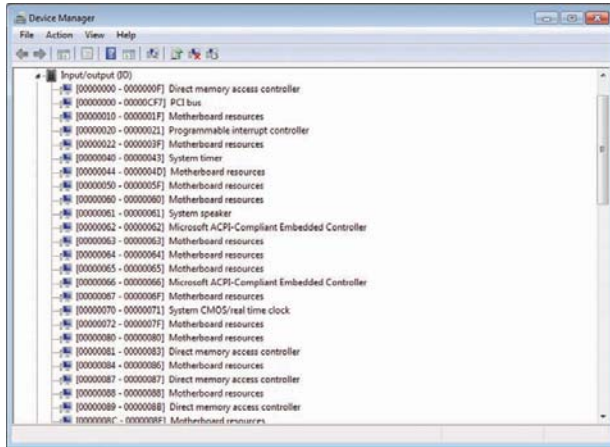
Input/Output (I/O) Port

The *Input/Output (I/O) port* is a small amount of memory assigned to a device that temporarily holds small amounts of data. It is used to transfer data between two locations. The data remains in the I/O port assignment until it can be moved.

The system has many hardware devices and software programs that depend on transferring data between hardware and memory locations. When a device or software program wants to transfer data, it may need to wait until the CPU is finished with its current process. In fact, it may need to wait until several processes are completed before the CPU can handle the data. The I/O port stores

Input/Output (I/O) port
a small amount of memory assigned to a device that temporarily holds small amounts of data. It is used to transfer data between two locations.

Figure 9-17.
I/O port
assignments.



this data until the CPU is free. Figure 9-17 shows the I/O port assignments associated with a system's hardware devices.

Memory Address Assignment

A *memory address assignment* is a large block of memory assigned to a device and is used to transfer data between two locations. A device is assigned a range of memory addresses. Some devices such as video and sound cards require a great deal of memory. The blocks of RAM assigned to a device cannot be used by any other device. Figure 9-18 shows the memory assignments associated with a system's hardware devices.

RAID Systems

Redundant Array of Independent Disks (RAID) is a system of disks arranged for speed or fault tolerance, or both. *Fault tolerance* as applied to RAID systems means the ability to recover from a hard disk or hard disk controller failure without the loss of stored data. Fault tolerance as applied to a network infrastructure means the ability to continue operation during a system hardware or software error. To achieve a high-data transfer rate, a technique known as *disk striping* is used. *Disk striping* involves dividing the data into separate sections and writing the data across several hard disk drives at the same time. This reduces the total amount of time it takes to store large amounts of data.

Tech Tip

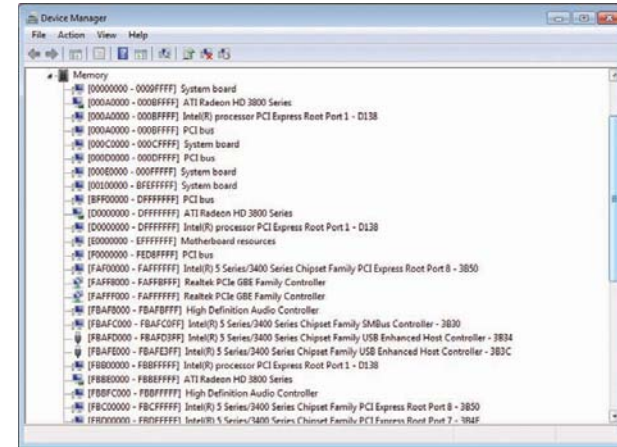


Another accepted representation of the RAID acronym is Redundant Array of Inexpensive Disks.

memory address assignment
a large block of memory assigned to a device. It is used to transfer data between two locations.

Redundant Array of Independent Disks (RAID)
a system of disks arranged for speed or fault tolerance, or both.

Figure 9-18.
Memory
assignments.



Another RAID technique is *error correction*, which can use traditional error-checking code (ECC) or parity. ECC is obsolete as a method for RAID because most hard disk drive systems use ECC as a standard way to protect data. Incorporating ECC into a RAID system would be redundant.

Parity is a technique that allows data to be recovered if one of the hard disk drives fails in a multiple disk drive system. A minimum of three hard disk drives must be used for this technique. If three hard disk drives are used, one hard disk drive is used to store parity and the other two are used to store data. Data is spread evenly between the two data storage drives, and the parity code for the sum of the two data storage drives is stored on the parity drive.

Parity is also a binary code that represents the total data pattern shared between the data storage drives. If any data storage drive fails, the system can use the parity bit to rebuild the missing data.

RAID systems are referred to as *RAID levels*. There are many RAID levels, each incorporating a different technique to increase speed or to provide data redundancy, or both.

Things to consider when selecting a RAID level are cost, reliability, and speed. There is no single, best RAID level. The appropriate level depends on the value of the data being protected and the amount of budget available. If a bank were protecting financial data, cost would most likely be of little concern. A RAID 0/5 with a tape backup would be quite appropriate. If you were protecting a personal computer, you would not likely need a RAID system.

Repairing a failed RAID system means the drive must be physically replaced and then the data reconstructed. In a Microsoft Windows operating system, the Disk Management utility is used to reconstruct data. If using proprietary SCSI drives, data is reconstructed by the data array manufacturer software package.

fault tolerance
as applied to RAID systems, it is the ability to recover from a hard disk or hard disk controller failure without the loss of stored data. As applied to a network infrastructure, it is the ability to continue operation during a system hardware or software error.

disk striping
a storage technique that divides the data into sections and writes the data across several hard disk drives at the same time.

error correction
a RAID technique that uses traditional error-checking code (ECC) or parity.



Tech Tip When implementing a RAID system in a Microsoft Windows operating system, all drives must be dynamic disk.

See **Figure 9-19** for a list of RAID levels and a short description of each. The following section looks at each RAID level in detail.

RAID 0

RAID 0 uses disk striping across a group of independent hard disk drives, **Figure 9-20**. This technique increases data flow but provides no fault tolerance. If one hard disk drive fails, all data is lost and cannot be reconstructed. This RAID level is most useful when speed is important, not fault tolerance.

RAID 1

RAID 1 uses the technique of disk mirroring. *Disk mirroring* is the act of writing the same information to two hard disk drives at the same time, **Figure 9-21**. Each of the two hard disk drives contains the same data. If one hard disk drive fails, a copy exists on the other hard disk drive.

Another form of RAID 1 is duplexing. *Duplexing* is the technique of placing each mirrored hard drive on a separate hard disk drive controller. Duplexing adds another level of fault tolerance. If one hard disk drive controller fails, the other is still operable. Mirroring contains the risk of the hard disk drive controller failing and causing both hard disk drives to fail. RAID 1 provides data protection at the cost of speed when compared to RAID 0. When there is a small demand for disk read/writes, RAID 1 is an appropriate technique.

disk mirroring
the act of writing the same information to two hard disk drives at the same time.

duplexing
the technique of placing each mirrored hard drive on a separate hard disk drive controller.

Figure 9-19. RAID levels.

RAID Level	Description
RAID 0	Striping.
RAID 1	Mirroring or duplexing.
RAID 2	Error checking code (ECC).
RAID 3	Byte-level striping with parity.
RAID 4	Block-level striping with parity.
RAID 5	Block-level striping with distributed parity.
RAID 6	Dual parity.
RAID 0/1 or 10	Disk mirroring with striping.
RAID 0/5 or 50	Block striping with parity and striping.

Figure 9-20. RAID 0 divides data and writes it to multiple drives. It provides excellent data transfer rates but no data loss protection.





Figure 9-21. RAID 1, or disk mirroring, provides a copy of data on two disk drives.

RAID 2

RAID 2 uses common ECC error-correction code to provide fault tolerance. All drives provide some sort of ECC, so this form of RAID is basically obsolete.

RAID 3

RAID 3 stripes data across multiple drives and stores parity on a separate drive. Data striping is performed at the byte level. See **Figure 9-22**.

RAID 4

RAID 4 stripes data across several drives and stores parity on a separate drive. The main difference between RAID 3 and RAID 4 is RAID 4 stores data in blocks. The size of the block can vary.

RAID 5

RAID 5 is also called *block striping with distributed parity*. RAID 5 distributes parity across all drives rather than writing parity to one drive. It provides fault tolerance and some increase in read-write data transfer. See **Figure 9-23**.




Figure 9-22. RAID 3 writes across multiple drives with parity stored on a separate drive.




Figure 9-23. RAID 5 provides both disk striping and parity. Parity and data is distributed across all drives.

RAID 0/1

RAID 0/1, or RAID 10, uses multiple RAID 1 (mirroring) disk systems. It also incorporates disk striping across each of the mirrored disk sets. When compared to RAID 1, this system provides better fault tolerance because data is duplicated on multiple mirrored disk sets. It also provides relatively fast data-transfer speeds by striping data across the mirrored disk sets.

RAID 0/5

RAID 0/5, or RAID 50, uses several RAID 5 (block striping with parity) sets and combines them with RAID 0 (disk striping). It provides increased fault tolerance when compared to a RAID 5 system.

RAID 6

RAID 6 is similar to RAID 5 in that it distributes parity across all drives, except RAID 6 uses a second set of parity. This technique is referred to as *dual parity*. The additional parity set allows for protection from simultaneous two-disk failure. Since two simultaneous disk failures are rarely encountered, RAID 6 is not commonly used.

**Network+ Note**

Pay particular attention to RAID 0, 1, and 5. These systems are the most commonly used and are the ones most likely to be covered on the Network+ Certification exam.

External Storage Systems

Data is often stored separately from a server. There are network devices that serve specifically as storage containers for network data. These devices are typically called *network attached storage (NAS)*. There are also facilities external to the local area network that provide data storage. These devices are typically called *storage area network (SAN)*. This section discusses both types of storage systems.

Network-Attached Storage (NAS)

Network attached storage (NAS) is a device or collection of devices that provide storage for network data. Network attached storage units are typically composed of disk arrays or tape arrays. The file systems used on the disk array and file array do not typically match the file system used by network clients or servers. The file systems are proprietary and are made to facilitate access speed and data integrity.

This is how data on a NAS device is accessed. A client requesting data makes a request to the file server. The file server in turn makes a request to the NAS device. The NAS device retrieves the data and sends it to the file server. The file server, in turn, sends the data to the client. See **Figure 9-24**.

Storage Area Network (SAN)

A *storage area network (SAN)* is a separate, high-speed network that provides a storage facility for one or more networks. Typically, a SAN uses a high-speed access media such as Fibre Channel. See **Figure 9-25**. Fibre Channel is discussed in the next section of the chapter.

network attached storage (NAS) a device or collection of devices that provide storage for network data.

storage area network (SAN) a separate, high-speed network that provides a storage facility for one or more networks.




Figure 9-24. In a network attached storage (NAS) design, clients make requests for data through the network file server. The file server then makes a request to the NAS device. The NAS device retrieves the data and sends it to the file server. The file server forwards the data to the client.




Figure 9-25. The storage area network (SAN) is a separate, high-speed network that provides a storage facility for other networks.

Exact definitions of storage area network vary according to different manufacturers. Often the differences are blurred.

Tech Tip

Fibre Channel a high-speed access method that typically uses fiber-optic cable as network media.

Network Address Authority (NAA) a naming standard that consists of the "naa." prefix followed by a string of hexadecimal characters. It was developed by the International Committee for Information Technology Standards (ICITS).

Fibre Channel

Fibre Channel is a high-speed access method that typically uses fiber-optic cable as network media. It may also use copper core cable and wireless. The term *Fibre Channel* also refers to a set of standards and a protocol. The Fibre Channel standard provides a relatively high-data transmission rate between supercomputers, mainframes, servers, and desktops. Fibre Channel is often used for storage area network (SAN) access. There are three common Fibre Channel topologies: point-to-point, arbitrated loop, and fabric, **Figure 9-26**.

A fiber channel device uses a **Network Address Authority (NAA)**, which is a naming standard developed by the International Committee for Information Technology Standards (ICITS). The NAA address consists of the "naa." prefix followed by a string of hexadecimal characters. For example, naa.02034D12B12124.

Fibre Channel Point-to-Point Topology

A Fibre Channel point-to-point topology is simply a straight connection or channel between two points, such as a mainframe and a server. The channel is isolated from other channels or media, and total bandwidth is dedicated to the channel between the two devices.

Fibre Channel Arbitrated Loop Topology

The Fibre Channel arbitrated loop is a common Fibre Channel topology. It looks similar to the Token Ring topology. The arbitrated loop differs from Token Ring in that two devices in the loop set up a direct communication link, or channel, for the duration of the data transfer. In Token Ring network, a token is passed around the ring while one device on the ring controls communication. Token Ring has a maximum control time of approximately 8 milliseconds in contrast to the arbitrated loop standard, which allows two devices to communicate until all the data requested has been transferred.

The arbitrated loop is limited to 127 connections. It is used to move large volumes of data between two devices such as a disk storage array and a server. It is not used for sporadic communication between multiple nodes as it is on a Token Ring network.

Fibre Channel Fabric Switched Topology

The fabric switched topology uses a device known as a *fabric switch*. A *fabric switch* is a switch designed especially for Fibre Channel networking. The fabric switch provides a direct or switched connection between two points. When two points wish to communicate, a private link, or channel, is set up between the two devices. This method is best used when numerous nodes wish to access the same disk storage array. For example, a university consisting of many buildings can connect each building's server to a common disk storage array using the fabric switched topology. This way, departments in each building can have access to research documents stored in the disk storage array. They can also access servers in other buildings through a high-speed connection.

fabric switch a switch designed especially for Fibre Channel networking and that provides a direct or switched connection between two points.

Internet Small Computer Systems Interface

Internet Small Computer Systems Interface (iSCSI) is a network storage standard developed by the Internet Engineering Task Force (IETF). It is used to access a wide variety of storage devices such as tapes, disks, and storage arrays. The iSCSI (pronounced eye-scuzy) standard is referred to as *IP-based storage* because it is designed to work on a LAN, MAN, or WAN using IPv4 and IPv6 network addresses to identify the storage device. iSCSI operates by sending SCSI commands over Ethernet using NFS or SMB/CIFS protocols. The two advantages of iSCSI are the following:

- Storage devices can be located at great distances from the client. This overcomes the limitations of systems such as Fibre Channel.
- It is less expensive than Fibre Channel SANs.

Windows 7 and Windows Server 2008 have a GUI interface called *iSCSI Initiator* for configuring iSCSI. It can be accessed from the **Administrative Tools** menu, **Figure 9-27**.

The iSCSI initiator is the workstation or server making the requests to the storage device. It is called the *target* and can be located either on the local area network or at some distant location across the Internet.

The iSCSI properties can be configured automatically through the Network Discovery feature or manually. Look at **Figure 9-28** to see how the Network Discovery feature can automatically locate an iSCSI device on the local area network. In this example, the iSCSI storage device is located on a Windows Server 2008 computer using IP address 192.168.1.103.

iSCSI devices can be located across the Internet by using an iSCSI qualified name (IQN) or an Enterprise Unique Name (EUN). The name can be automatically generated based on information provided by the Network Discovery feature.

Figure 9-29 shows an example of an IQN address. The IQN address in the example is iqn.1991-05.com.microsoft:richar7.netlab1.domain.local. The IQN name format is <type>-<date>-<authority>-<string defined by naming authority>:

where

- Type = iqn
- Date = month and year domain was acquired
- Authority = domain name
- String defined by naming authority = String of characters identifying the iSCSI object

Internet Small Computer Systems Interface (iSCSI) an IP-based storage technology that uses IPv4 and IPv6 network addresses to identify storage devices on a LAN, MAN, or WAN.

Figure 9-26. Three major topologies associated with Fibre Channel systems are point-to-point, arbitrated loop, and fabric switched.

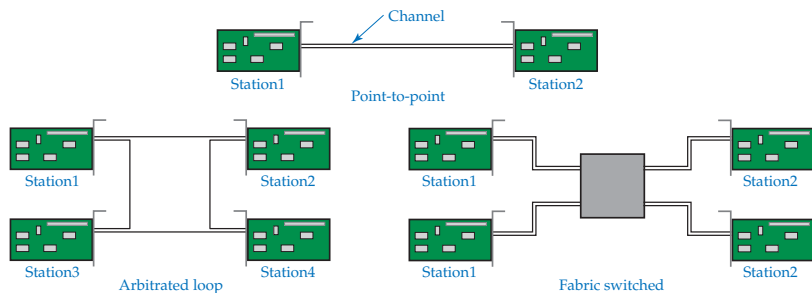


Figure 9-27.
The iSCSI Initiator, located under the **Administrative Tools** menu, is a GUI used for configuring an iSCSI system.




Figure 9-28.
Discovery tab of the **iSCSI Initiator Properties** dialog box.



Figure 9-29.
Configuration tab of the **iSCSI Initiator Properties** dialog box.



The string of characters consists of letters A–Z, numbers 0–9, and special symbols dash (-), period (.), and colon (:). The string of characters typically provides a general description, but this is not a requirement. The string of characters can represent anything. Everything before the colon is a mandatory standard. Everything after the colon is entirely up to the naming authority.

The IQN address in the example is first identified by “iqn”. Following this are the date of origination of the address “1991-05,” the reverse domain name of the authority “com.microsoft,” and the computer name and domain.

The IEEE has developed another naming format known as the *Extended Unique Identifier* or the *IEU-64* format. The IEU format consists of the prefix “ieu” followed by a period and then by a string of hexadecimal numbers. For example, ieu.0123D0123F0A. This type of address is issued by IEEE and is used by equipment manufacturers to identify its network storage devices.

Furthering Your Study

You may be interested in further study and possible certification as a server technician. The CompTIA organization has a certification called *Server+*, which is designed to verify your competence for installing, repairing, and maintaining servers. You can download a complete description of the exam objectives from the CompTIA Web site.

Many of the areas in the Server+ Certification exam overlap those on the A+ Certification exam. However, even if you have scored high on the A+ Certification exams, there are sufficient differences between the certifications that will prevent you from passing the Server+ Certification exam without specialized study.

For additional study materials and more detailed information about server hardware, access the IBM and the Sun Microsystems Web sites. For detailed information on SCSI technology, access the Adaptec Web site.

Summary

- Servers play many roles in a network, such as print server, file server, database server, application server, backup server, Web server, and mail server.
- A thin server has only the hardware and software needed to support and run a specific function, or role.
- A thin client server provides applications and processing power to a thin client.
- The thin client relies on a thin client server's processor and memory for processing data and running software applications.
- Typically, entry-level servers contain 1 to 8 processors; mid-range servers contain 9 to 30 processors; and high-end servers contain 31 to 106 processors.
- Servers typically incorporate hot-swap technology to allow components to be removed or installed while the system is running.
- System resources associated with hardware are Direct Memory Access (DMA), interrupt request (IRQ), I/O port address, and memory address assignment.
- When two hardware devices share the same resource assignment, a system conflict occurs.
- The most common RAID levels are RAID 0, RAID 1, and RAID 5.
- RAID 0 is known as striping. Striping increases data read/write speed but does not provide fault tolerance.
- RAID 1 is mirroring or duplexing two hard disk drives.
- RAID 5 is known as striping with parity, which combines data read/write speed with fault tolerance.
- Network Attached Storage (NAS) is a device or collection of devices used to provide storage for a local area network. The NAS shares the bandwidth of the local network.
- Storage Area Network (SAN) is a separate network dedicated to data storage. It does not affect the bandwidth of the other connected networks.
- The iSCSI standard is an IP-based storage technology that uses IPv4 and IPv6 network addresses to identify storage devices on a LAN, MAN, or WAN.

Review Questions

Answer the following questions on a separate sheet of paper. Please do not write in this book.

1. What are some typical services that servers provide?
2. What is a thin server and what applications can it be used for?
3. What is a thin client server and what applications can it be used for?
4. What does the term *hot-swap* mean?
5. What type of motherboard is commonly used for servers?
6. Which server component is responsible for the POST?
7. How many devices can be connected to a Fast/Wide SCSI-2 chain?
8. What is the bus width of Fast SCSI-2?
9. What is the bus width of Fast/Wide SCSI-2?

10. What must be done to the last device on a SCSI chain?
11. What does the acronym LUN represent?
12. What does the acronym ACPI represent?
13. What is the purpose of ACPI?
14. What does the acronym ACPICA represent?
15. What is the goal of ACPICA?
16. Name the four system resources that may be assigned to a hardware device?
17. How many IRQ settings are there in a typical computer?
18. What does the acronym RAID represent?
19. How is parity used to replace data lost on a damaged disk?
20. Which RAID level(s) provide no fault tolerance?
21. Which RAID levels provide disk mirroring?
22. What is the difference between disk mirroring and duplexing?
23. Which RAID level provides both striping and parity?
24. Your system storage is accessed many times throughout the day by many users. The data is not critical. Rather, it is a collection of application software. A tape backup system is used to restore data from a system failure. Which RAID level would *most likely* be appropriate and why?
25. Which RAID level is denoted by the use of two hard disk drives each containing identical data?
26. What does the acronym NAS represent?
27. What does the acronym SAN represent?
28. What is the difference between NAS and SAN?
29. What type of media is used for Fibre Channel?
30. What are the three common Fibre Channel topologies?
31. What does the acronym iSCSI represent?
32. What are the two main advantages to iSCSI storage?
33. What two protocols are used to send SCSI commands across the network?
34. What does the acronym NAA represent?
35. What two naming formats are used to identify iSCSI targets?
36. What naming format is used to identify Fibre Channel storage devices?
37. Identify the following address type: iqn.2008-06.com.rmroberts:myiSCSIstorage.

Sample Network+ Exam Questions

Answer the following questions on a separate sheet of paper. Please do not write in this book.

1. Which RAID level provides the fastest reading and writing data transfer time?
 - A. RAID 0
 - B. RAID 1
 - C. RAID 5
 - D. RAID 0 and RAID 1
2. Which network RAID level provides a duplicate copy of one hard disk drive on another hard disk drive?
 - A. RAID 0
 - B. RAID 1
 - C. RAID 5
 - D. RAID 4



Network+



3. How does SAN differ from NAS?
 - A. SAN provides access to a SCSI disk exclusively, while NAS provides access to SCSI, ATA, and EIDE drives.
 - B. SAN offers encryption; NAS does not.
 - C. SAN does not increase local network traffic; NAS does.
 - D. SAN is connected to the local network system, while NAS is accessed remotely via a WAN link.
4. What is the minimum number of drives required for implementing a RAID 5 installation?
 - A. 1
 - B. 2
 - C. 3
 - D. 4
5. What does the term *fault tolerance* mean?
 - A. The ability to find faults in data before it is backed up.
 - B. The ability of the network administrator to tolerate system policy breaches by users.
 - C. The ability to recover from a system failure.
 - D. The ability to continue normal operation despite a system failure.
6. Which of the following describes the ability to change a module such as hard disk drive or power supply on a server without the need to shut down the server?
 - A. System substitution
 - B. On the fly changes
 - C. Hot-swap
 - D. Infinite run ability
7. What is the function of disk striping?
 - A. Reduces the total amount of time it takes to store large amounts of data.
 - B. Stores duplicate data across two or more drives.
 - C. Ensures data integrity.
 - D. Aids in disaster recovery.
8. How does a client access files on a NAS device?
 - A. Through NAS client software.
 - B. Through a server that is attached to the NAS device.
 - C. Directly through Category 5e cable.
 - D. Directly through high-speed, fiber-optic cable.
9. An investment broker is installing a network system for his employees and to provide customer access. Money transactions will occur on a continuous basis. Which RAID level would you recommend for this business?
 - A. RAID 0
 - B. RAID 1
 - C. RAID 1 (with duplexed drives)
 - D. RAID 5
10. Multiple CPUs in a server can perform which of the following?
 - A. Unilateral processing
 - B. Multiplexed processing
 - C. Synchronous processing
 - D. Parallel processing

Suggested Laboratory Activities

Do not attempt any suggested laboratory activities without your instructor's permission. Certain activities can render the PC operating system inoperable.

1. Create several partitions on a hard disk drive and then format the partitions using FAT32.
2. Convert a basic FAT32 partition to an NTFS partition and then convert the NTFS partition to dynamic disk.
3. Use the **diskpart** command to view the contents of a hard disk drive. Familiarize yourself with the DiskPart command line options.
4. Install a RAID 0, RAID 1, and RAID 5 system. Remove a hard disk drive to simulate a disk failure and then replace the hard disk drive with a new hard disk drive and rebuild the missing data.
5. Explore the iSCSI Initiator administrative tool in Windows 7.

Interesting Web Sites for More Information

<http://computer.howstuffworks.com/scsi.htm>

[http://msdn.microsoft.com/en-us/library/ee681827\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ee681827(VS.85).aspx)

<http://serverwatch.internet.com>

www.acnc.com/raid.html

www.faqs.org/rfcs/rfc3720.html

www.molex.com

www.pcguides.com/topic.html

www.starwindsoftware.com/white-papers

An enterprise business typically depends on hundreds of servers to meet user demands.



Chapter 9 Laboratory Activity

Using the DiskPart Command Interpreter

After completing this laboratory activity, you will be able to:

- Use the DiskPart command interpreter.
- Use common DiskPart commands to retrieve information about the number of disks installed in a system and about disk partitions or volumes.
- Use DiskPart **help** command to reveal a list of DiskPart commands.

Introduction

Disk management is usually performed in the Disk Management utility. The DiskPart command interpreter is an alternative way to manage disks and partitions. At times, the DiskPart command interpreter may be the only alternative for inspecting and manipulating the master boot record (MBR) and partition table should the GUI fail. Also, certain tasks can be performed using the DiskPart command interpreter that cannot be performed using the Disk Management utility. For example, you must use the DiskPart command interpreter to convert a dynamic disk to a basic disk and to delete or replace the MBR in the active partition. The Disk Management utility does not allow the performance of these tasks because it is designed to protect the integrity of the disk structure from accidental deletion.

You can think of the DiskPart command interpreter as an updated version of the **fdisk** command. Windows XP no longer supports the **fdisk** command. The DiskPart utility is used in place of **fdisk** to inspect, create, and delete partitions or volumes on a disk drive.

The DiskPart command interpreter will not make changes to removable media such as a USB drive and an IEEE-1394 drive. These types of media are identified as a “super floppy” by the DiskPart utility rather than as a hard disk drive.

In this laboratory activity, you will perform some basic DiskPart operations to become familiar with this utility. While you most likely will never need to use the full capability of this utility, more information can be obtained at the Microsoft Tech Support Web site.

Look at the following screen capture. Notice that the **detail disk** command has been entered at the DiskPart command line. The **detail disk** command reveals details about the selected disk, such as the hard disk drive manufacturer and type. It also shows that the hard disk drive has two partitions or volumes. Volume 1 is assigned letter *C* and volume 2 is assigned letter *D*. Notice that both volumes are indicated as healthy.


```

Command Prompt - diskpart
DISKPART> detail disk
TOSHIBA MK3018GAS
Disk ID: 6013B998
Type: IDE
Bus: 0
Target: 0
LUN ID: 0

Volume ### Ltr Label Fs Type Size Status Info
-----
Volume 1 C NTFS Partition 16 GB Healthy System
Volume 2 D NTFS Partition 12 GB Healthy

DISKPART>

```

The **list** command can be used with the disk, volume, and partition object. When using the DiskPart command interpreter, disks, partitions, and volumes are referred to as *objects*. Look at the following screen capture. It shows the **list disk**, **list partition**, and **list volume** commands. The **list disk** command lists the disk(s) installed starting with disk 0. The **list volume** command lists the volumes contained in the system. The **list partition** command lists the partitions located on a selected disk. Notice how the DVD-ROM drive is identified after issuing the **list volume** command. Also, notice how the removable drive is identified even though it is no longer present.

```

Command Prompt - diskpart
DISKPART> list disk
Disk ### Status Size Free Dyn Gpt
---
Disk 0 Online 28 GB 0 B

DISKPART> list partition
Partition ### Type Size Offset
-----
Partition 1 Primary 16 GB 32 KB
Partition 2 Extended 12 GB 16 GB
Partition 3 Logical 12 GB 16 GB

DISKPART> list volume
Volume ### Ltr Label Fs Type Size Status Info
-----
Volume 0 E DVD-ROM 0 B
Volume 1 C NTFS Partition 16 GB Healthy System
Volume 2 D NTFS Partition 12 GB Healthy
Volume 3 F Removable 0 B

DISKPART>

```

Before issuing the **list partition** command, you must select the disk you wish to view. For example, to view the partition information for a single hard disk drive, you would issue the **select disk** command. Then you would issue the **list partition** command. Otherwise, you will generate an error message stating, "There is no disk selected to list partitions."

Selecting the disk is also known as *setting focus*. When you set focus, you focus the **list partition** command on a selected disk. You must also set the focus on the partition you wish to list. If you do not, the default partition you are presently running the DiskPart command interpreter from is automatically selected.

For more information on the DiskPart command interpreter, read the following Microsoft Support article: 300415—A Description of the DiskPart Command Line Utility.

Equipment and Materials

- Computer with Windows XP, Windows Vista, or Windows 7 installed. (A server is not required for this laboratory activity.)

Procedure

- Report to your assigned workstation.
- Boot the computer and verify it is in working order.
- Open the command prompt (**Start | All Programs | Accessories | Command Prompt**, or **Start | Run** and enter **cmd** in the **Run** dialog box).
- At the command line, type and enter the **DiskPart** command. If successful, the prompt will display **> DISKPART** to let you know that you are now using the DiskPart command interpreter.
- Enter the **Help** or **/?** command to reveal a list of commands.
- Enter the **list** command to display the objects that can be specified with the **list** command. Disk, partition, and volume are listed.
- Try each of the following commands to reveal information about the disk drive(s): **list disk**, **list partition**, **list volume**. Remember, the **list partition** command requires focus. You must use the **select** command to select the disk drive you wish to list partition information about. For example, you would enter **select disk 0** to set the focus on the first disk in the system. The **select** command is only required for the **list partition** command. After running each command, record on a separate sheet of paper the information revealed.
- The **detail** command reveals similar content but with more detailed information about an object. Enter each of the following commands and then record on a separate sheet of paper the information revealed: **detail disk**, **detail volume**, **detail partition**.
- Practice using the commands covered thus far until you are very familiar with them. They are difficult to learn at first because some involve setting focus and using two words for a valid command.
- Use the **help** command to reveal what the following commands are used for. Do *not* attempt to use any of the following commands without explicit permission from the instructor.

Command	Purpose
add	
active	
assign	
break	
clean	
create	
delete	
detail	
exit	
extend	
help	
import	
inactive	
list	
online	
rem	
remove	
repair	
rescan	
retain	
select	

11. _____ Enter the **exit** command to return to the command prompt and then again to return to the Windows desktop. Leave your computer in the condition as specified by your instructor.
12. _____ Answer the review questions.

Review Questions

Answer the following questions on a separate sheet of paper. Please do not write in this book.

1. What are objects when using the DiskPart command interpreter?
2. What command is used to reveal information about a particular partition?
3. Which command is used to set focus on a particular drive?
4. What does the command prompt look like when you are running the DiskPart command interpreter?
5. How can you see a list of the various DiskPart commands?
6. A particular computer has only one hard disk drive. What command is used to set focus on that one hard disk drive?
7. What command will stop the DiskPart command interpreter and return you to the default prompt?