

CISCO

IT Essentials I

Chapter 9

Advanced Hardware and Servers



Chapter 9 Study Guide

Advanced Hardware and Servers

Please keep this study guide safe. Each student will be given one copy of the study guide. They will be needed each lesson when we are studying this unit. Replacement copies will be issued at a charge of 50p per copy.

Name: _____

Network Servers

A network server is a computer system in a network that is shared by multiple users. Also, one keyboard, monitor, and mouse may connect to any number of servers through a keyboard video mouse (KVM) switch. KVM switch is a device used to connect one keyboard, one mouse, and one monitor to two or more computers. KVM switches are used to save space on a desktop when two or more computers are routinely used. They are also widely used to control servers that are only accessed periodically. Switches often use special cables that combine keyboard, monitor, and mouse cables into one port at the switch end.




RAID—Redundant Array of Inexpensive Disks

RAID is designed to allow some fault tolerance to prevent loss of data in the event of a disk drive failure on a network server. A disk drive is a mechanical device, which might eventually fail. RAID accomplishes this fault tolerance or redundancy by storing the information on more than one disk drive.

The most common RAID levels that we are going to use are:

- RAID 0**
- RAID 1**
- RAID 5**
- RAID 0/1**



RAID 0

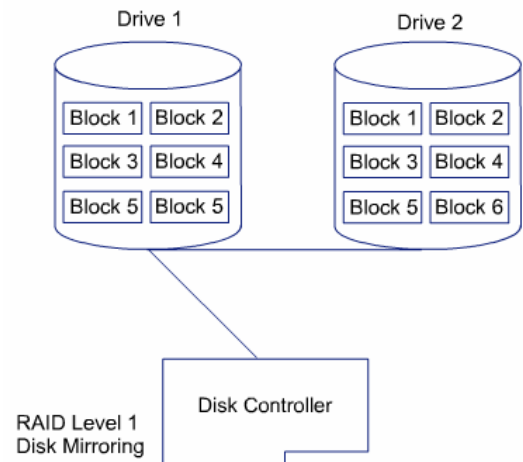
RAID 0 is known as disk striping. Specifically, it is a stripe set without parity. RAID 0 is not fault tolerant, but it is used to improve disk input/output performance. RAID 0 should not be used in a production-server environment. However, RAID 0 is often used in a high-powered workstation to improve disk input/output performance by reading and writing files in blocks to several disks simultaneously as opposed to reading and writing a file sequentially to a single disk drive. To implement RAID 0, at least two disk drives are needed. For example, three 18-GB disk drives configured to implement RAID 0 have a storage capacity of 54 GB.

RAID 1

RAID 1 has two different implementations. They are disk mirroring and disk duplexing.

Mirroring

In disk mirroring, everything written to one disk is also written to a second disk. Reading can be performed from either disk. Since data is duplicated on two different disk drives, the failure of one disk drive is not a serious problem, but merely an inconvenience. The network server will keep running using the single working disk drive. When time permits, the broken disk drive can be replaced, and the data on the working disk drive can be copied to the replacement. This recreates the mirror. A minimum of two disk drives is required to implement RAID 1. In disk mirroring, both disk drives are attached to the same disk controller.

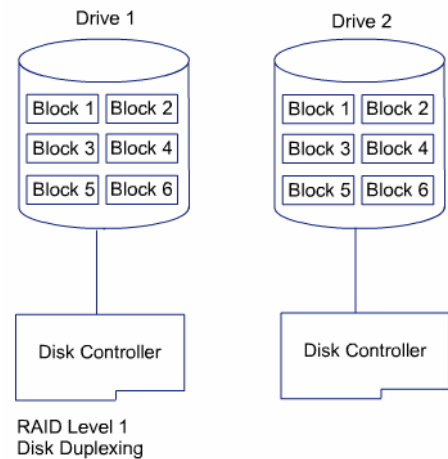


RAID 1—Mirroring


Duplexing

RAID 1 Configuration (Duplexing)

Disk duplexing eliminates the single point of failure that exists in disk mirroring. This is done by adding another disk controller and configuring the RAID system to duplicate data on disk drives that are attached to two different disk controllers. There is generally no significant performance difference between disk mirroring and disk duplexing. The user is just adding further redundancy in the form of a second controller.



RAID 1—Duplexing

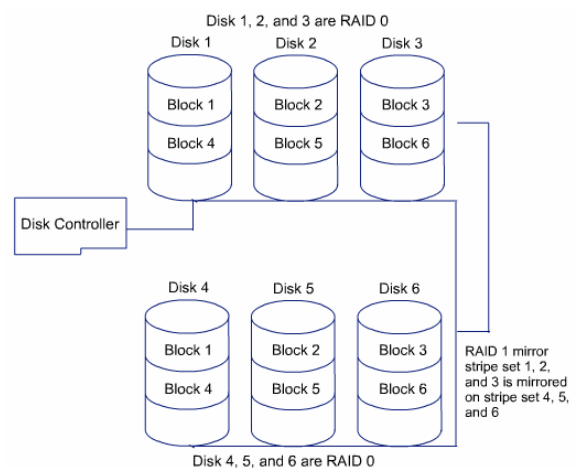


RAID 5

RAID 5 uses a much more complicated scheme to provide fault tolerance in the case of a single disk failure. RAID 5 requires a minimum of three disk drives to implement. The disk drives that comprise a RAID 5 solution are often referred to as a RAID 5 array. The failure of a single disk drive does not cause the network server to fail. The missing information that was on the failed disk can be recreated quickly using the information on the remaining disks. The failed disk drive should be replaced as quickly as possible. RAID 5 cannot survive the failure of a second disk drive after one disk drive has failed. Some RAID systems allow for the configuration of a "hot spare" disk drive in the RAID system. A hot spare disk drive is powered up and running, but it contains no data. It is just waiting for a drive in the disk array to fail so that it can be used.

RAID 0/1

RAID 0/1, which is sometimes called RAID 0+1 or RAID 10, involves mirroring or duplexing two RAID 0 arrays. This yields the fault tolerance of RAID 1 and the input/output speed of RAID 0. RAID 0/1 requires a minimum of four disk drives to implement. Here is an example of a RAID 0/1 implementation.



Reflection

1. How would multiple servers connect to one keyboard mouse and monitor?
2. List the four common RAID configurations:
3. Which RAID configuration is usually used for high powered workstations and not in a server environment?
4. **True or False:** RAID 1 requires at least two hard disks:
5. What is the difference between RAID 1 Mirroring and RAID 1 Duplexing?
6. Which RAID configuration provides fault tolerance, spreads the information among each disk and allows for one disk in the array to fail without losing any data?
7. Which RAID level provides the performance of RAID 0 along with the redundancy of RAID 1?

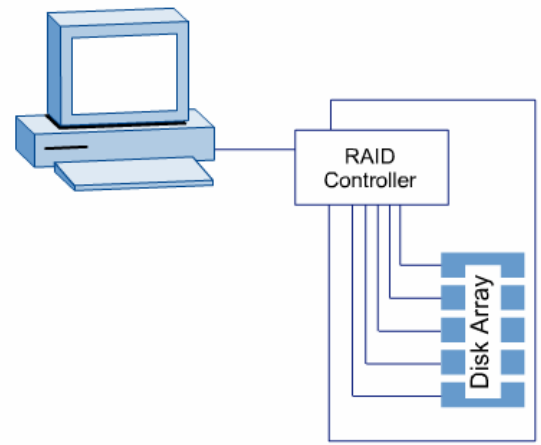
Tip: Learn the five most commonly used RAID configurations.

RAID Controllers

RAID controllers are specialized disk controllers that use either advanced technology attachment (ATA) or small computer system interface (SCSI) technologies. ATA uses a maximum of two channels with a maximum of two disk drives per channel for a total of four disk drives. SCSI RAID controllers have multiple channels. Two channels are common. RAID controllers with three, four, and five channels are available. RAID controllers are generally expensive due to the sophistication that they must contain.

RAID controllers often have an onboard memory cache ranging in size from 4 MB to 256 MB. This onboard memory cache often has a battery backup system to prevent data loss in the event of sudden power loss to the network server.

The memory cache on the RAID controller can usually be configured as read cache, write cache, or a combination of both. The read cache will improve the read performance. The write cache will allow the processor to continue with other tasks instead of waiting for the data to be written to the disk.



Hardware Versus Software RAID

RAID is usually implemented using a RAID disk controller. However, RAID disk controllers are rather expensive. RAID can also be implemented in software by several network operating systems, including Novell NetWare, Linux Read Hat, Microsoft Windows NT, and Microsoft Windows 2000.

When using the Windows 2000 version of RAID, the hard drive must be converted to a dynamic disk before the RAID options are available to implement.

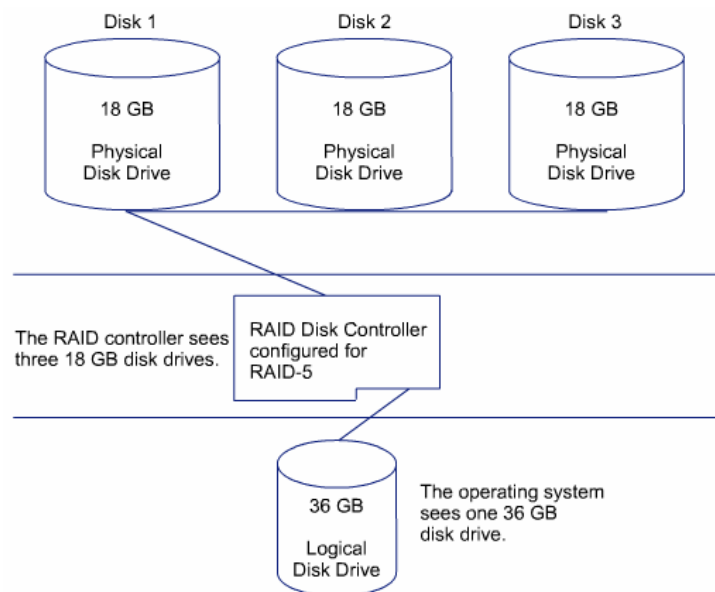
Software RAID systems usually support RAID 0, 1, and 5. Software RAID is usually implemented at the disk partition level rather than the physical disk as in hardware RAID.

When RAID 5 is implemented in software, the files on the RAID array are not available until the network server operating system is running.

Configuring Hardware RAID Controllers

Network servers that contain a RAID controller must have the RAID system configured before the network operating system can be installed. Configuration of the RAID system consists of selecting actual physical disk drives and grouping them together into one of the available RAID configurations. These configurations are usually RAID 1 or RAID 5.

RAID is used to provide fault tolerance in case of a disk drive failure in the network server. The term hardware-based RAID means that the disk drives in the network server have RAID implemented by a special disk controller, which is the RAID controller.



The formula to work out the amount of storage on a RAID 5 array is $(n - 1) * c$, where n is the number of disk drives and c is the capacity of each of the disk drives.

Question?

If a RAID 5 array consists of 6 x 10GB hard disks, what would the storage capacity be?

Reflection

1. Usually in a workstation Hard disks connect to an IDE controller on the motherboard. When using a RAID setup what are the two types of controller that can be used? (IDE is not one of them)

2. How many channels can each of these two types of controller use?

3. RAID controllers have onboard memory cache that is backed up by an onboard battery in case of power loss. What is the advantage of having write cache on a RAID controller?

4. What is the main difference between software RAID and hardware RAID?

5. **True or False:** When using a hardware RAID system the RAID configuration must be configured before installing the network operating system.

6. When setting up a software RAID configuration the hard disk must be converted to a _____ disk.

***Tips:* Know what type of RAID controllers can be used.**

Know the difference between hardware and software RAID configurations.

Know what needs to be done to any hard disks before setting up software RAID

External Peripherals

External disk subsystems are necessary when the amount of disk storage cannot be accommodated by the disk drive bays internal to the network server chassis. These external disk subsystems can be either SCSI or Fibre Channel. External CD-ROM systems are generally used to implement CD-ROM libraries, which can accommodate a large number of CD-ROM drives and make them available to client computers on the network. The network servers that implement CD-ROM libraries are often called CD-ROM servers. Even though server class microcomputers often have many empty bays designed to hold disk drives, it is often necessary to have disk drives external to the actual server chassis. External disk subsystems may consist of a single disk drive in its own chassis with its own power supply. On the high end, an external disk subsystem chassis might have 100 or more disk drives in it.



CDROM Servers

Attaching that many CD-ROM drives to a single network server is a fairly simple process. It is done through a little-used feature called a Logical Unit Number (LUN). Although LUNs are defined in the SCSI standards, they are seldom used except on large groups of CD-ROM drives. A LUN enables a user to assign sub-SCSI IDs to a single SCSI ID. This means that the user could have 7 CD-ROM drives all with the SCSI ID of 5. They could each have a different LUN of 1 through 7 all on the same SCSI channel. This means that on a single SCSI channel with SCSI IDs of 1 through 7, each SCSI ID could have 7 LUNs. This would make a total of 49 CD-ROM drives on a single SCSI channel.

Upgrading Network Servers With Additional Processors

To add another processor to a multiprocessor-capable network server, the new processor must meet the following criteria:

- Be the same model processor, which includes Pentium, Pentium Pro, Pentium II, Pentium II Xeon, Pentium III, Pentium III Xeon, Pentium 4, and so on, as the existing processor
- Have the same clock speed
- Have the same Level 2 (L2) cache size
- Match the stepping within one version (N+1)

Windows NT Server 4.0

If the original installation of Microsoft Windows NT Server 4 was on a network server with a single processor, the hardware abstraction layer (HAL) on the network server must be updated for the network server to recognize and use multiple processors. To upgrade Windows NT 4 to a multiprocessor HAL, use the UPTOMP.EXE utility available on the Microsoft Windows NT 4 Server Resource Kit.

Windows 2000 Server

If the original installation of Microsoft Windows 2000 Server was on a network server with a single processor, the HAL on the network server must be updated for the network server to recognize and use multiple processors. To install support for multiple processors on Windows 2000, follow the procedure in Step-by-Step 4.3, which is from Microsoft Knowledge Base Article Q234558.

Novell NetWare 5

To upgrade Novell NetWare 5 so that it will recognize the additional processor, follow these steps:

- Step 1 Load NWCONFIG | Multi CPU Options | Select a Platform Support Module.
- Step 2 Restart the Novell NetWare 5 network server after NWCONFIG modifies the STARTUP.NCF and the AUTOEXEC.NCF files.

Adding Hard Drives

Disk drive upgrades come in two varieties. The first type of upgrade involves adding disk drives to an existing network server, and the second type involves replacing existing disk drives with larger or faster disk drives. Upgrades to disk drives have the most potential of any upgrade to destroy data. Before attempting any disk drive upgrade, make sure that there is at least one, preferably two, verified full backups of the data on the disk drives.



ATA Hard Disk

The ATA channels are usually labelled primary and secondary so that the system can distinguish between them. When only a single disk drive is attached to the ATA disk controller, a second disk drive can be attached in either of the two following ways:


- To the primary controller using the same ribbon cable as the first hard disk
- To the secondary controller using a different ribbon cable

This is very similar to the way in which IDE hard drives connect to computers.

Putting the second ATA disk drive onto the secondary channel results in having one ATA disk drive on each ATA channel. Then performance of the disk subsystem can be enhanced.

Adding SCSI Hard Disk Drives

Adding SCSI disk drives to an existing SCSI channel is a rather simple process. See Figure . To make sure the addition works, the server administrator or hardware specialist must review the documentation of the SCSI bus. They need to know the SCSI IDs of the existing disk drives and where the SCSI bus is terminated. For internal SCSI devices, they also need to determine whether any SCSI connectors are available on the SCSI bus ribbon cable. See Figure . If no SCSI connectors are available on the SCSI bus ribbon cable, obtain a new SCSI ribbon cable with the correct number of connectors. Then set the SCSI ID of the new disk drive to a SCSI ID that is not already in use on the SCSI bus. The SCSI terminator might also need to be removed. Remember that the SCSI bus must be terminated at both ends.



SCSI Hard Disk

External SCSI devices are usually connected to the SCSI channel in a daisy-chain manner. Adding an additional external SCSI device involves picking an SCSI ID that is not currently in use on the SCSI channel and adding the SCSI device into the daisy chain. Exceeding the SCSI cable length is the biggest problem encountered when adding an additional external SCSI device. The second most common problem is proper termination of both ends of the SCSI bus.

Adding SCSI Drives to a RAID Array

Adding drives to an SCSI-based RAID array is no different from adding drives to an SCSI channel. The only exception is that after the disk drives are added to the array, RAID configuration utility must be used to add the disk drives to the RAID array.

Remember to be able to boot an operating system from a SCSI drive, the SCSI drive must have a SCSI ID of 0.

A SCSI system has a memory buffer which takes care of all the read write operations. This means the computers processor is free to perform other tasks.

Reflection

1. External disk subsystems use either _____ or _____

2. **True or False:** When attaching many CD-ROM drives to a single network server you can assign 7 sub-SCSI IDs to one single SCSI ID.

3. When upgrading to a dual processor setup in Novell NetWare 5, which program must be run in order for the operating system to recognize the additional processor?

4. After adding a second processor to a Windows NT server, the PC technician discovers that the operating system is only using the original processor. What must the technician do to get the second processor running?
 - a) follow the procedure in Step-by-Step 4.3, which is from Microsoft Knowledge Base Article Q234558.
 - b) run the UPTOMP.EXE utility.
 - c) Load NWCONFIG | Multi CPU Options | Select a Platform Support Module.

5. Which is the best way to configure two ATA disk drives? (Remember ATA uses two channels)
 - a) Both disks on the primary ATA channel with one ribbon cable.
 - b) One disk on each ATA channel
 - c) Both disks on the secondary ATA channel with one ribbon cable.

6. In order to boot from a SCSI hard disk the hard disk must have a SCSI ID of _____.

7. A SCSI chain must be terminated at?
 - a) Start of the chain
 - b) Middle of the chain
 - c) Both ends of the chain
 - d) End of the chain

8. In a SCSI system all the read and write operations occur on the_____.

Know how external disk subsystems connect to a server.

Understand SCSI sub ID's.

Know what needs to be done to add additional processors to an Netware and NT server.

Understand how ATA drives connect to a computer.

Understand how SCSI drives connect to a computer.

Adding Memory


It has been said that there is no such thing as too much memory in a server. Although in many cases this is true, there are a few exceptions. One, the user can only put as much memory in the network server as it was designed to contain. There is always a maximum amount of memory that can be supported by the processors and/or the control chipsets of the motherboard of the network server. The other exception is having more memory than the network server operating system can utilize. Keep both of these exceptions in mind when considering a memory upgrade to a network server.

Before adding memory to a network server, the user needs to verify the current memory configuration. The most reliable way to check the existing memory configuration is to open the chassis of the network server. Try to answer the following questions:


- How many memory slots does the network server have?
- How many memory slots are empty and available for additional memory to be installed?
- What is the size in MB and speed of the current memory modules?
- What type of memory module is currently installed? Memory module types include single in-line memory modules (SIMMs), dual in-line memory modules (DIMMs), RAMbus in-line memory module (RIMMs), buffered, unbuffered, registered, and so on.
- What type of memory is the memory module that is used in the network server? Memory types include extended data out (EDO), dynamic random access memory (DRAM), synchronous DRAM (SDRAM), or Rambus DRAM (RDRAM).
- What error-detection method is used on the memory module? Methods include parity, non-parity, ECC, and non-ECC.
- Answer these questions by reviewing the documentation that was shipped with the network server. This information might also be found in the log that was kept as part of the installation process.

Installing Additional Memory—SIMM's and DIMM's

After determining the feasibility of the memory upgrade and the compatibility of the memory with the network server, the last step is to actually install the additional memory. Be sure to install the memory according to the manufacturer instructions. Remember that the memory modules may have a requirement to be installed in pairs or groups of four. Most network servers have SIMMs, DIMMs, or RIMMs.



SIMM Installation



DIMM Installation

Installing Additional Memory—RIMM's

RIMM installation differs only slightly from DIMM installation. All memory module slots designed to use RIMMs must be populated. If an actual RIMM memory module is not installed in a slot, a continuity module must be installed. A continuity module does not contain any memory. Adding a RIMM memory module involves removing a continuity module and replacing it with a RIMM memory module. Failure to have continuity modules in the memory slots not occupied by RIMMs may result in a network server that does not power up.

Upgrading SCSI Adapter

Upgrading the BIOS or firmware on an adapter is very similar to upgrading the system BIOS. The steps are vendor specific, but the general steps are shown below:

- Step 1. Locate the latest BIOS or firmware on the adapter vendor's website.
- Step 2. Download the BIOS or firmware upgrade and follow the vendor's instructions to install the upgrade.

Replacing a PCI Adapter

Adapters are generally replaced after they fail. The replacement procedure generally follows several simple steps, as shown in Figure):

This process requires that the network server be taken out of operation. This results in downtime and lost productivity. However, a recent technology known as peripheral component interconnect (PCI) hot plug or PCI hot swap enables the user to replace, upgrade, or add an adapter without powering down the network server. PCI hot plug has three capabilities:

Hot replacement – Removing a failed PCI adapter and inserting an identical adapter into the same slot while the network server is operational.

Hot upgrade – Replacing an existing adapter with an upgraded adapter while the network server is running.

Hot expansion – Installing a new adapter into a previously empty slot while the network server is running.

For PCI hot plug to work, the network server hardware, the adapter drivers, and the network server operating system must be PCI hot plug aware. The network server hardware allows power to be removed from individual PCI slots and allows adapters to be removed and inserted without the use of a screwdriver. On a PCI bus that supports hot plug, a slot release lever replaces the use of a screw to secure the adapter in the PCI slot.

Peripheral Devices

A peripheral device is any device that is not part of the core computer system, which includes the processor, memory, and the data bus. Peripheral devices can be either internal to the server chassis or external to the network server chassis. Internal peripherals include such components as disk drives, CD-ROM drives, floppy disk drives, and network interface cards. External peripherals are devices external to the network server chassis, such as printers, modems, monitors, keyboards, and mice. External devices such as printers and monitors might be hot-swappable. Other external devices might require that the network server be shut down before they can be upgraded.



Reflection

1. How much memory can be put into a computer?

2. All memory module slots designed to use RIMMs memory must be _____.

3. Where is the best place to download the latest BIOS upgrade for an adapter?

4. _____ enables a user to replace, upgrade, or add an adapter without powering down the network server.

5. What is a peripheral device?

Know how RIMM modules are used in computers.

Know where to get upgrades for devices.

Know what a peripheral device is.