Module 9: Advanced Hardware Fundamentals and Servers

Module Overview

Overview

Module 9: Advanced Hardware and Servers

Upon completion of this module, students will be able to complete tasks related to the following:

- 9.1 Network Server Overview
- 9.2 Hardware-Based RAID Configuration
- 9.3 Configuring External Peripherals
- 9.4 Adding Hardware to a Server
- 9.5 Upgrading Server Components

A network server is the center of a network environment. The server allows users to access files, e-mail, programs, and printers. Fault tolerance is important for a network server, because it allows a system to continue when a hardware failure occurs. One method used to provide fault tolerance is the Redundant Array of Inexpensive Disks (RAID) technology. This module focuses on RAID and discusses memory upgrades, the configuration of external disk subsystems, and external compact disk read-only memory (CD-ROM) systems.

9.1 Network Server Overview

9.1.1 Network server

A network server is a computer system in a network that is shared by multiple users. Figures 1 to 3 are examples of different network servers. Servers come in all sizes from x86-based PCs to IBM mainframes. A server may have a keyboard, monitor, and mouse directly attached. Also, one keyboard, monitor, and mouse may connect to any number of servers through a keyboard video mouse (KVM) switch, which is shown in Figure 4. Servers may also be accessed through a network connection.

NOTE:
A 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

The term server may refer to both the hardware and software, which is the entire computer system, or it may refer to just the software that performs the service. For example, the e-mail server may refer to the e-mail server software in a system that also runs other applications. It may also refer to a computer dedicated only to the e-mail server application.

9.1.2 RAID

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.

RAID level 1 uses duplication of the data to provide fault tolerance. RAID levels 3, 4, and 5 use parity information that is calculated from the bit patterns of the data being written to the RAID array to provide fault tolerance. When a disk drive fails in RAID 3, 4, or 5, the parity information can be used along with the data on the remaining disk drives in the array to calculate the data that was on the disk drive that failed. This allows the disk subsystem and the network server to keep functioning. However, it will be slightly slower because of the calculations required to re-create the missing data. RAID level 2 is structurally different in that it does not use duplication or parity to provide fault tolerance. RAID 2 uses a special hamming code instead.

RAID is a term that is surrounded by a tremendous amount of misinformation. There is disagreement about how many levels of RAID are defined, whether the A in RAID stands for array or arrays, and whether the I in RAID stands for inexpensive or independent. In the past few years, many people have substituted the word independent for inexpensive.

RAID 6, 7, 10, 50, 53, and others can be found in the literature provided by many vendors. This module will focus on the types of RAID used most often in a network environment.

RAID was defined in 1987 in the paper called, "A Case for Redundant Arrays of Inexpensive Disks (RAID)", which was written by David A. Patterson, Garth A. Gibson, and Randy H. Katz at the University of California, Berkeley. The original paper defined five levels of RAID and offered the RAID solution as an alternative to single large expensive disk (SLED).

RAID 0
RAID 0 was not defined in the 1987 Berkeley paper. In fact, it is not RAID because it does not provide any redundancy. RAID 0 is just an array or group of disk drives used as a single disk. The data is written in chunks or stripes to all the disk drives in the array. This improves disk input and output performance because several chunks of data can be written or read simultaneously. If a disk drive in the RAID 0 array fails, all data in the RAID 0 array is lost. RAID level 0 is also often called disk striping without parity. Figure 1 shows an illustration of RAID 0.

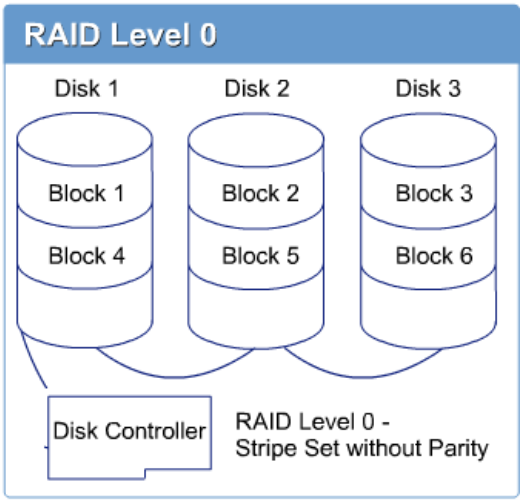

RAID 1
RAID 1 requires a minimum of two disk drives. All other RAID levels, except level 0, require at least three disk drives to implement. RAID 1 writes all data to two separate locations. To store 20 gigabytes (GB) of data using RAID 1, two 20-GB disk drives are required. This is a 50 percent loss of storage capacity.

There are two ways to implement RAID 1:

- Disk mirroring
- Disk duplexing

In disk mirroring, the two disk drives are connected to the same disk controller. The only problem with disk mirroring is that if the disk controller fails, there is no access to the mirrored data. Figure 2 shows a diagram of disk mirroring. To eliminate this single point of failure, use disk duplexing rather than disk mirroring.

In disk duplexing, each disk drive in the mirrored set is connected to a different disk controller. This eliminates the single point of failure in pure disk mirroring. The only additional cost is the additional disk controller. Figure 3 shows a diagram of disk duplexing.

RAID 2
RAID 2 uses a hamming code to create an error correcting code (ECC) for all data to be stored on the RAID 2 array. The ECC can detect and correct single-bit errors and detect double-bit errors. The ECC code has to be read and decoded each time data is read from the disk. RAID 2 is very difficult and expensive to implement and has a very high overhead. For example, there are three parity bits for each four data bits.

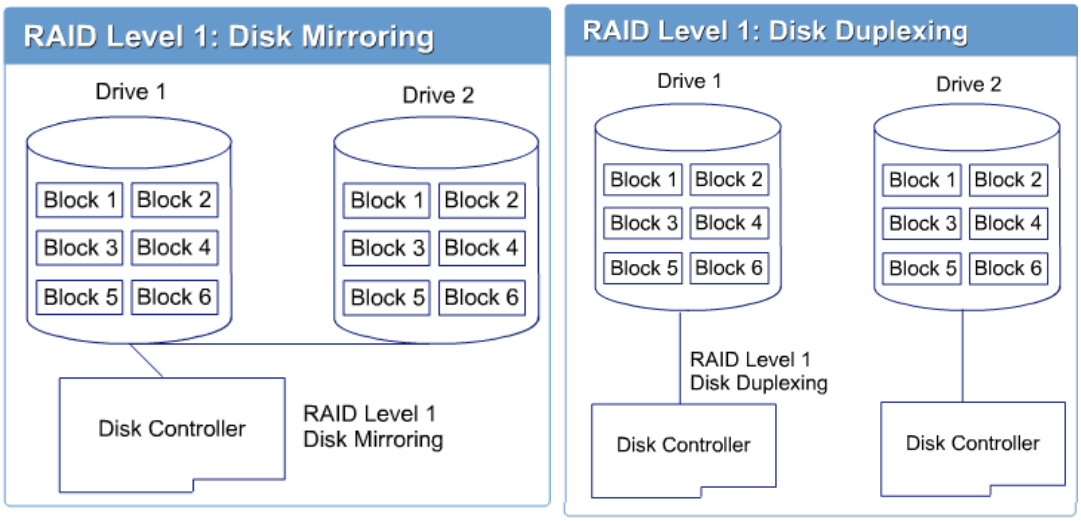

NOTE:

A hamming code is an error correction method that mixes three check bits at the end of each four data bits. When these check bits are received, they are used to detect and correct one-bit errors automatically.
RAID 2 has no commercial implementations because of the expense and difficulty of implementation. It requires a minimum of three disk drives to implement.

RAID 3

RAID 3 uses bit-level parity with a single-parity disk to provide fault tolerance of data stored on the RAID 3 array in the event of failure of a single disk drive in the array. RAID 3 requires that all the disk drives in the array be synchronized with each other. The bits of the data and the parity information calculated from the data are written to all the disk drives in the array simultaneously. RAID 3 requires a minimum of three disk drives to create the array.

RAID 4

RAID 4 uses block-level parity with a single-parity disk to provide fault tolerance to the RAID 4 array in the event of failure of a single disk drive in the array. On a RAID 4 array, data and the parity information calculated from the data is written to the disk drives in blocks. There is no need for the disk drives to be synchronized together, and the disk drives can be accessed independently. A minimum of three disk drives is required to create the array. The problem with RAID 4 is that the parity drive is accessed on every write operation to the RAID array. This will cause heavy utilization of the parity drive, which will probably fail before the other drives in the array.

RAID 5

RAID 5 uses block-level parity, but it spreads the parity information among all the disk drives in the disk array. This eliminates the parity drive failure common in RAID 4 systems. The loss of storage capacity in RAID 5 systems is equivalent to the storage capacity of one of the disk drives. If there are three 10-GB disk drives in a RAID 5 array, the storage capacity of the array will be 20 GB, which is a loss of one-third, or 33 percent. In another example, if there are seven 10-GB disk drives in a RAID 5 array, the total storage capacity of the array will be 60 GB, which is a loss of one-sixth, or 16.67 percent. Figure 4 shows a diagram of RAID 5.

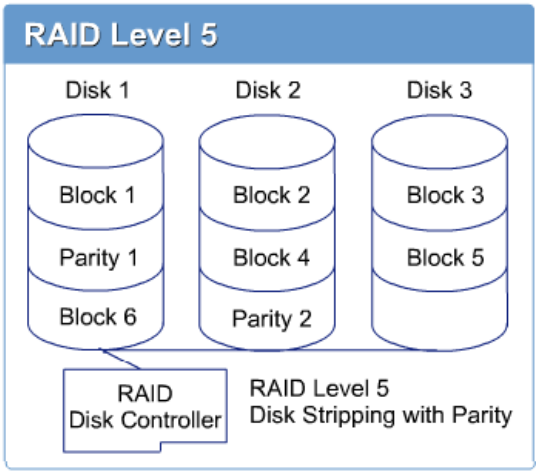

RAID 0/1

RAID 0/1 is also known as RAID 0+1, and it is sometimes called RAID 10. This combination of RAIDs provides the best of both worlds. It has the performance of RAID 0 and the redundancy of RAID 1. RAID 0/1 requires at least four disk drives to implement. In RAID 0/1, there are two RAID 0 stripe sets, which are used to provide high input/output performance, that are mirrored. This provides the fault tolerance. Figure 5 shows a diagram of RAID 0/1.

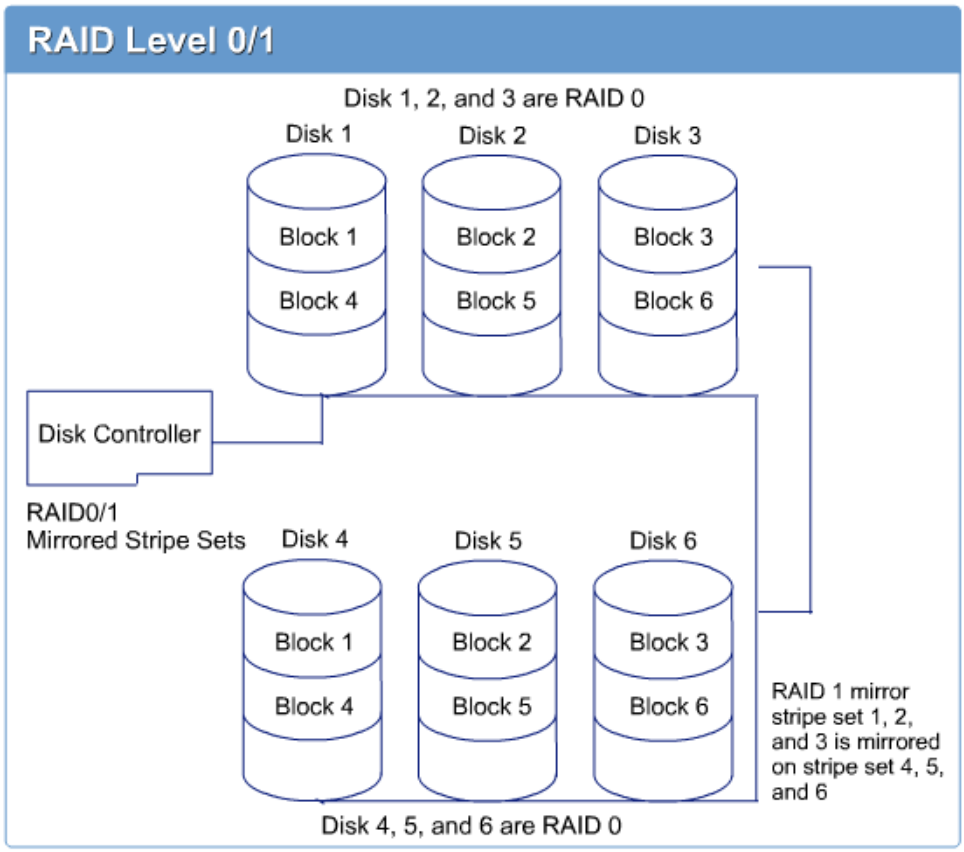

9.1.3 RAID controller

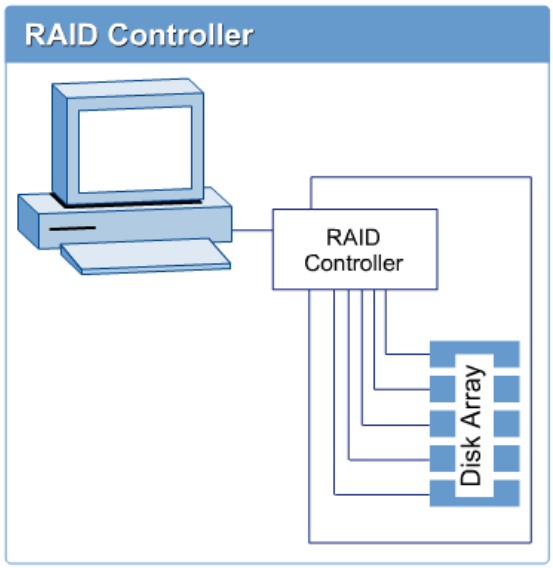

RAID controllers are specialized disk controllers that use either advanced technology attachment (ATA) or small computer system interface (SCSI) technologies. ATA RAID controllers are limited in the number of disks that can be attached. This is due to ATA channel limitations, which are a maximum of two 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. This is important because data written from the system memory to the RAID controller is first written to the onboard cache, and it could be several seconds before the data is actually written to the disk. Data on the disk drive may not be updated with the current data without the battery to supply power to the RAID controller. This could easily lead to the loss of data integrity.

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. Figure 1 shows the RAID controller and the disk arrays.

The following features should be considered when evaluating RAID controllers:

- Number of channels
- Speed of channels
- Onboard cache, which is read, write, combination, and battery backup option
- Fast host adapter PCI
- Bus width, which includes 16 bit, 32 bit, and 64 bit

9.1.4 Hardware RAID 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. See Figure 1.

NOTE:
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. The disadvantage to software RAID is that it requires the network server processor to perform the work usually done by the RAID controller in hardware RAID. RAID 5 implemented in software requires the processor calculate all the parity information when writing data to the RAID 5 disk array. RAID 1 implemented in software puts a minimal load on the network server processor.

When RAID 5 is implemented in software, the files on the RAID array are not available until the network server operating system is running. This means that the operating system cannot be stored on and therefore cannot boot from a RAID 5 system implemented in software. This is not an issue when a RAID 5 system is implemented using hardware.

Software-based RAID does have one advantage over hardware-based RAID. In software-based RAID, the RAID implementation can be based on disk partitions rather than entire disk drives. For example, three 10-GB partitions on three different disk drives can be used to create a software RAID 5 array. There could be space in different partitions on each of these disk drives that could be used for some other purpose. In nearly all cases, hardware-based RAID is better than software-based RAID. However, having RAID implemented in software is much better than not having any disk fault tolerance at all.

NOTE:
Understand the difference between hardware and software RAID.

9.2 Hardware-Based RAID Configuration

9.2.1 Hardware-based RAID configuration overview

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. The network hardware vendor or the RAID hardware vendor usually supplies software to aid in the configuration of the RAID system. The disk drives of the RAID system might be internal to the network server chassis or external in a separate enclosure.

As mentioned in the previous section, 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. Some network operating systems can implement software-based RAID at the expense of an additional load on the network server processor. Most RAID controllers are designed to use SCSI disk drives. However, at least one disk controller manufacturer makes a RAID controller that uses an EIDE/ATA-2 disk drive. The RAID disk controller has its own processor to implement the RAID configuration, and this relieves the network server processor of this task.

The configuration of the RAID controller in the network server is accomplished by software provided by the network server or RAID controller vendor. Although vendor specific, all the software works basically the same way. It enables a user to see the disk drives attached to the RAID controller. First, the user selects what disk drives to utilize. Then the user needs to specify the version of RAID to implement using the selected disk drives. The software then "prepares" the disk drives to implement the RAID solution. For example, the user might pick two physical disk drives and tell the RAID configuration software to use these two disk drives to implement a RAID 1 or mirroring solution. The RAID controller would tell the network server operating system that there is a single logical disk drive. In this case, the RAID controller is actually reading and writing to two physical disk drives.

In another example, a user might select five physical disk drives and tell the RAID configuration software to use these five physical disk drives to implement RAID 5, which is disk striping with parity. The RAID controller would tell the network server operating system that there is a single logical disk drive. In this case, the RAID controller is actually reading and writing data in blocks across all five disk drives in the RAID 5 disk array.

In yet another example, a user might select the same five physical disk drives and tell the RAID configuration software to use these five physical disk drives to implement RAID 5, which is disk striping with parity. In addition, the single logical disk drive could be partitioned into two partitions by the RAID configuration software. The network operating system would see two logical disk drives or one on each partition. In this case, the RAID controller is actually reading and writing data in blocks across all five disk drives in the RAID 5 disk array.

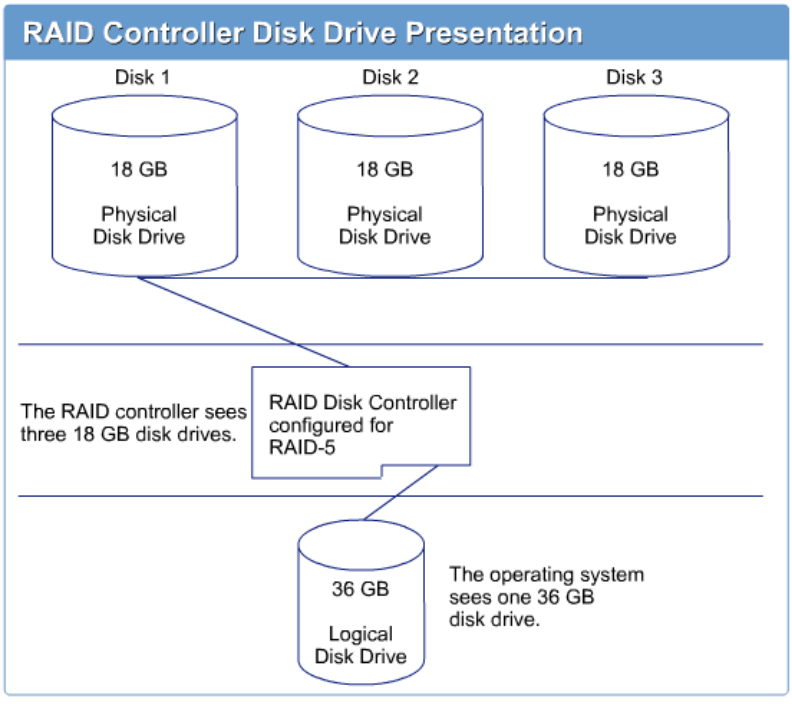

NOTE:

The Server+ exam will ask questions regarding general configurations and general understanding of RAID.

Figure 1 shows an example of how a RAID controller manages disk drives and shows the logical disk drive to the network server operating system.

9.2.2 RAID 0 configuration

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.

Figure 1 shows an example of a RAID 0 implementation. For example, two 18-GB disk drives configured to implement RAID 0 have a storage capacity of 36 GB.

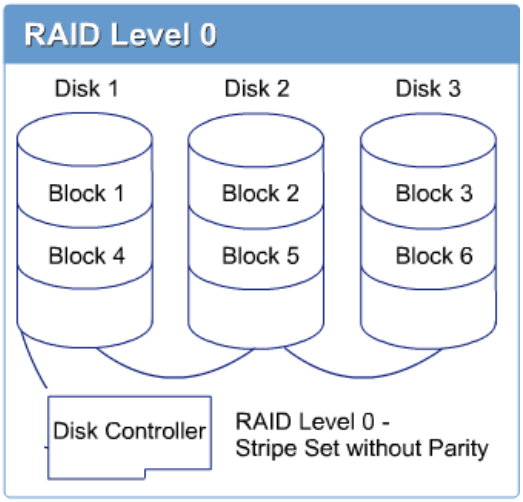

9.2.3 RAID 1 configuration

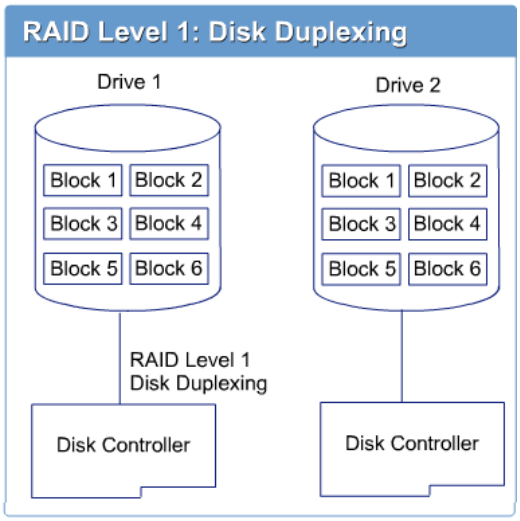

RAID 1 Configuration (Duplexing)
Disk duplexing eliminates the single point of failure that exists in disk mirroring. This is done by adding another disk controller and configuring the RAID system to duplicate data on disk drives that are attached to two different disk controllers. There is generally no significant performance difference between disk mirroring and disk duplexing. The user is just adding further redundancy in the form of a second controller. The overhead of RAID 1 duplexing is 50 percent. Figure 2 shows an example of a RAID 1 duplexing implementation.

9.2.4 RAID 5 configuration

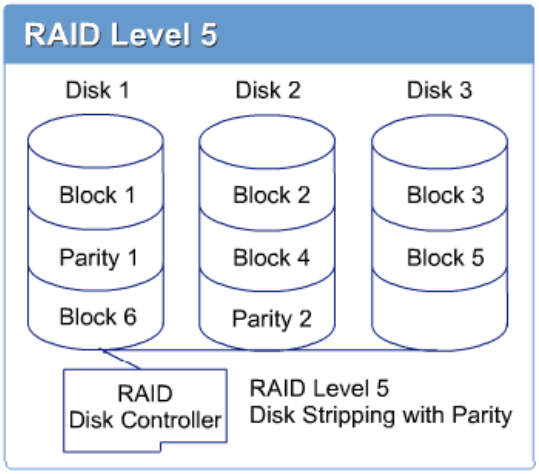

RAID 5 uses a much more complicated scheme to provide fault tolerance in the case of a single disk failure. Refer back to Section 9.1.2, RAID, for an in-depth discussion of RAID 5.

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. Because of this fact, 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.

When the failure of a disk drive in the RAID array occurs, the RAID system starts rebuilding the data that was on the failed drive on the hot spare disk drive. This hot spare methodology minimizes the amount of time it takes to get the RAID rebuilt. It also minimizes the window of time that the RAID system is vulnerable to a

second drive failure that could destroy all data stored on the RAID array. Figure 1 shows an example of a RAID 5 implementation.

RAID 5 is more efficient than the other RAID levels in that the overhead is 1/n * 100, where n is the number of disk drives in the RAID 5 array. In other words, if a RAID 5 array is comprised of six 18-GB disk drives, the overhead is 1/6 * 100, or 16.7 percent. Another way of looking at this is that the user loses the capacity of one of the disk drives in the RAID 5 array. This space is used to store the parity information. The parity information is actually stored across all the drives in the RAID 5 array.

The total storage capacity of the 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. In this example, the total storage capacity of the RAID 5 array is (6 – 1) * 18, or 90 GB.

9.2.5 RAID 0/1 configuration

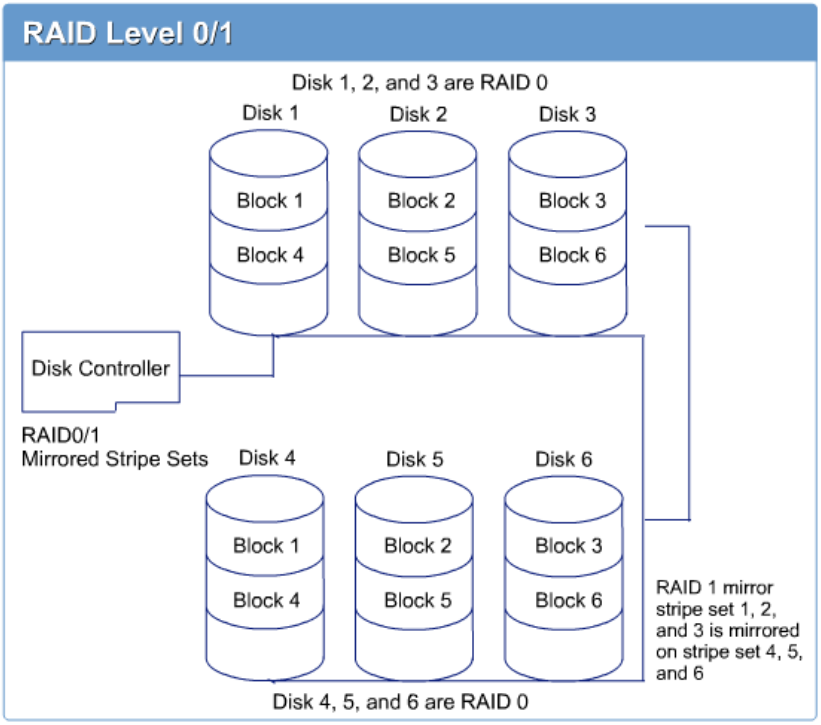

RAID 0/1, which is sometimes called RAID 0+1 or RAID 10, involves mirroring or duplexing two RAID 0 arrays. This yields the fault tolerance of RAID 1 and the input/output speed of RAID 0. RAID 0/1 requires a minimum of four disk drives to implement. Figure 1 shows an example of a RAID 0/1 implementation.

9.3 Configuring External Peripherals

9.3.1 Overview of external disk subsystems

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. Generally, Fibre Channel-based systems can support many more disk drives than a SCSI-based external system. 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. See Figure 1. The network servers that implement CD-ROM libraries are often called CD-ROM servers.

9.3.2 Configuring an external disk subsystem

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. See Figure 1.

The simple external disk subsystem with only a few disk drives might just be attached to the external port on an SCSI or RAID controller. The external disk drives would then function in the same way that internal disk drives function, except they are external to the network server chassis.

Some large external disk subsystems may have their own RAID mechanism built in. These large systems are often configured separately from the disk controller in the network server to which they will be attached.

Often a user can configure large external disk subsystems to be shared by more than one network server. This is one way to implement a high-availability server solution.

To connect the external disk subsystem to the network server, use a standard external SCSI cable or even Fibre Channel.

NOTE:

Fibre Channel can be configured point-to-point through a switched topology or in an arbitrated loop (FC-AL) with or without a hub, which can connect up to 127 nodes. It supports transmission rates up to 2.12 Gbps in each direction, and 4.25 Gbps is expected.

Generally, Fibre Channel-based external disk systems can handle a very large number of disk drives. In all cases, be sure that the power switch on the external disk subsystem is turned on before turning on the network server.

9.3.3 Configuring a external CD-ROM system

External CD-ROM systems are often referred to as a CD-ROM library, as shown in Figure 1. Imagine a tower chassis with 7, 14, 21, or more CD-ROM drives. Having these CD-ROM drives attached to a network server means that it is possible to share all these CD-ROM drives and the CD-ROMs that they contain with all users on the network.

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. To configure the external CD-ROM system, follow the manufacturer installation and configuration instructions. Be sure that the external CD-ROM system is powered up before powering up the network server.

NOTE:

Configuring external peripherals as in the external subsystems and CD-ROM systems is important to understand. It is necessary to understand why as well as how to configure external peripherals as in the external subsystems and CD-ROM systems.

9.4 Adding Hardware to a Server

9.4.1 Replacing a single processor with a faster processor

Figures 1 to 3 show the components of a server. Use those figures to reference to the components throughout the module.

Internal View

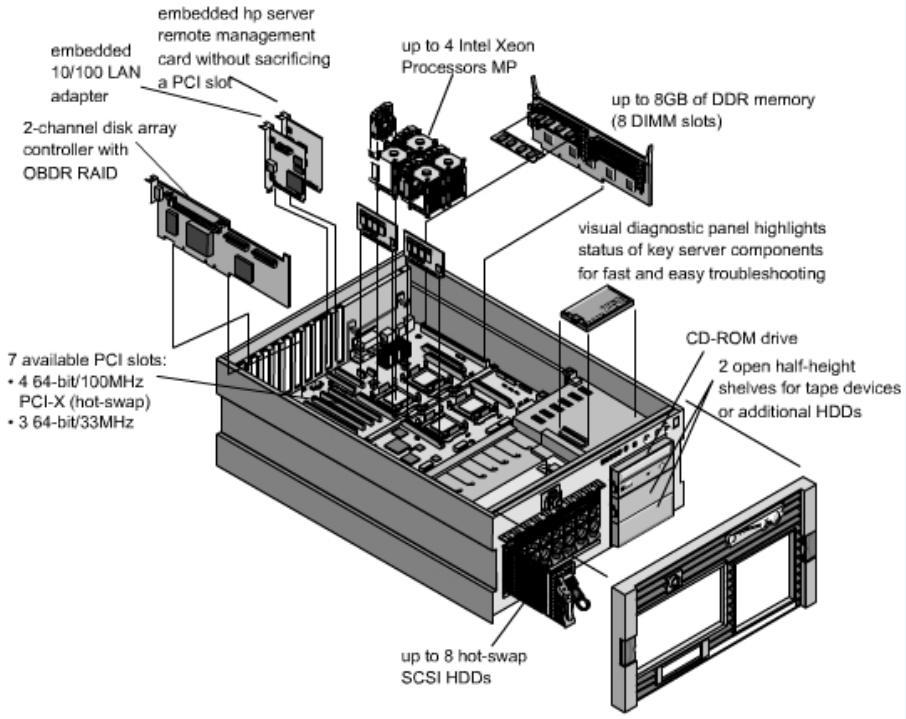

Deciding whether a processor in a network server can be replaced with a faster processor depends on several factors. The most important factor is whether the motherboard in the network server will support a processor with a faster clock cycle. The other factors include the physical package that the existing processor uses and whether a faster processor is available that utilizes the same physical package or form factor. Users can obtain this important upgrade information from the manufacturer of the network server motherboard. Check the motherboard manufacturer website to see whether the processor can be upgraded to a faster one. The user has to determine whether a faster processor is available in a form factor that is compatible with the existing processor. Upgrading to a faster processor may also require upgrading the BIOS on the system board. The following steps are needed in order to perform a single processor upgrade, as shown in Figure 4.

Hot Swapable Power Supply

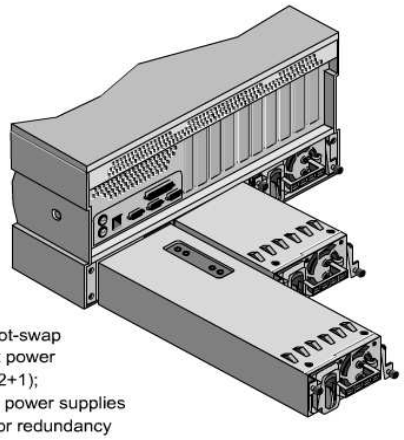

Side View

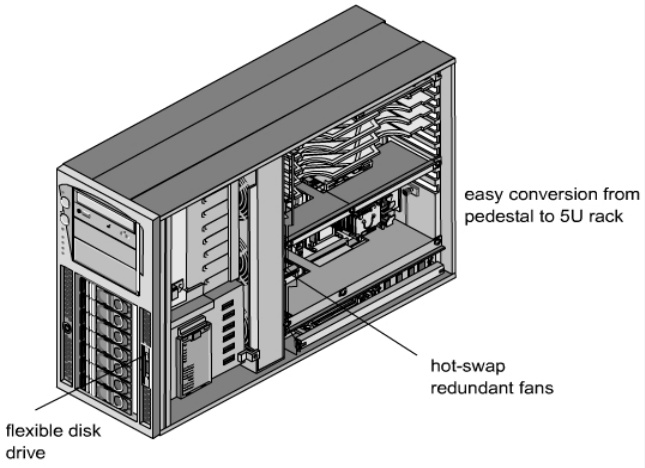

Replacing a Single Processor

- Step 1 - Follow the upgrade checklist.
- Step 2 - Upgrade the system BIOS.
- Step 3 - Open the network server chassis (following ESD best practices).
- Step 4 - Remove the current processor.
- Step 5 - Insert the new processor.
- Step 6 - Close the network server chassis.
- Step 7 - Verify that the new processor is recognized by the network server hardware and the network operating system.

9.4.2 Installing 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)

In order to tell what processor is currently in the network server so that the process can be matched, Intel provides information on all of its processors on its website. Intel also offers a utility that detects and identifies the Intel processor that is currently in a network server.

The Intel processor identification utility is available in two versions. One operates under the Microsoft Windows operating system, and one runs from a bootable DOS floppy disk.

The Microsoft Windows version of the Intel identification utility can be downloaded from http://support.intel.com/ support/ processors/ tools/ frequencyid/ freqid.htm.

Figure 1 shows sample output from the Windows version of the Intel identification program. The bootable floppy disk version of the Intel identification utility can be downloaded from http://support.intel.com/ support/ processors/ tools/ frequencyid/ bootable.htm.

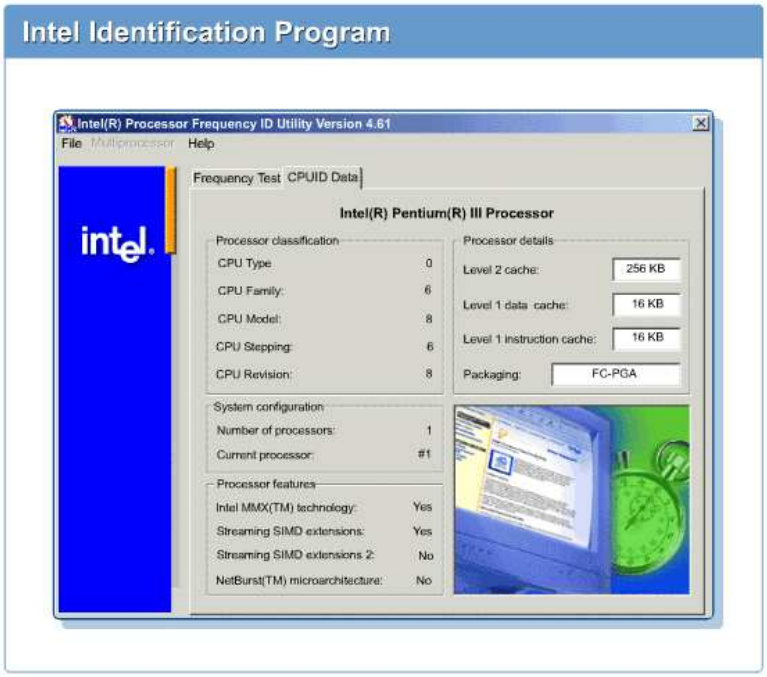

Figure 2 shows sample output from the bootable floppy disk version of the Intel identification program. A description of the Intel identification program along with installation instructions can be found at http://support.intel.com/ support/ processors/ tools/ frequencyid/ download.htm.

A tag on the processor itself can also identify the current processor. This tag contains a 5-digit S-spec number that starts with the letter S. The user can use this 5-digit number and the Intel website to identify the Intel processor that is in the network server.

Use the S-spec number from the processor label and the following Intel websites to identify the processor currently in the network server:

- Pentium http://developer.intel.com/ design/pentium/qit/
- Pentium Pro http://developer.intel.com/ design/pro/ qit/index.htm
- Pentium II http://support.intel.com/support/ processors/ sspec/p2p.htm
- Pentium II Xeon http://support.intel.com/support/ processors/sspec/p2xp.htm
- Pentium III http://support.intel.com/ support/ processors/ sspec/p3p.htm
- Pentium III Xeon http://support.intel.com/support/ processors/ sspec/ p3xp.htm
- Pentium 4 http://www.intel.com/support/ processors/ pentium4/

The processor to be added to the network server should come with installation instructions. The installation instructions for Intel processors are also available at the Intel website.

Intel processor installation manuals can be obtained at http://support.intel.com/ support/ processors/ manuals/.

The following steps are used when a adding an additional processor:

Step 1
Follow the upgrade checklist.

Step 2
Open the network server chassis by following ESD best practices.

Step 3
Insert the new processor.

Step 4
Close the network server chassis.

Step 5
Upgrade the system BIOS.

Step 6
Upgrade the operating system to recognize multiple processors.

Step 7
Verify that the new processor is recognized by the network server hardware and the network operating system.

Output from the Bootable Output Version

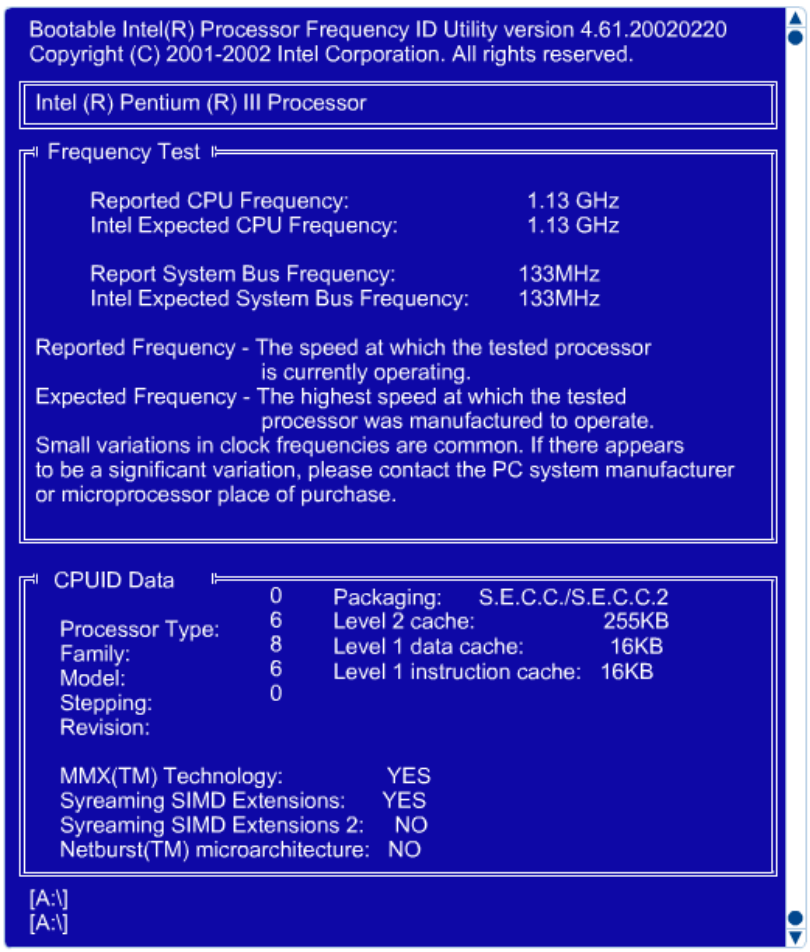

9.4.3 Upgrading the operating system for multiple processors

Quite often a network server operating system will not recognize that an additional processor has been installed on the network server. The following sections detail how to correct this situation.

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.

For information about the processors recognized by the current version of Microsoft Windows NT Server 4, open a command prompt window and type the following command:

set

This command prints a list of all the current environment variables. Look for the variable Number_of_processors to see how many processors Microsoft Windows NT Server 4 recognizes. The number of processors can be determined in other ways, but this is one of the easiest.

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.

Steps for Activating Multiple Processor Support on Windows 2000 Server are shown in Figure 1.

For information about the processors recognized by the current version of Microsoft Windows 2000 Server, open a command prompt window and type the following command:

set

This command prints a list of all the current environment variables. Look for the variable Number_of_processors to see how many processors Microsoft Windows 2000 Server recognizes. The number of processors can be determined in other ways, but this is one of the easiest.

Novell Netware 5
If the original installation of Novell NetWare 5 was on a network server with a single processor, several changes in the configuration of the Novell NetWare 5 server must be made for it to recognize and use multiple processors. 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.

For information about the processors recognized by the current version of Novell NetWare 5, use the following command on the NetWare console:

display processors

Red Hat Linux
For Red Hat Linux, and other versions of Linux, to recognize multiple processors, the Linux kernel must be rebuilt. Make sure that the main Makefile, which is /usr/src/linux/Makefile contains the line SMP=1. Rebuild the Linux kernel using the normal methods. For information about the processors recognized by the current version of Linux, use the following command:

cat /proc/cpuinfo

Activating Multiple Processor Support

- Step 1. Click **Start > Settings > Control Panel,** and then click **System.**
- Step 2. Click the **Hardware** tab, and select **Device Manager.**
- Step 3. Double-click to expand the **Computer** branch. Note the type of support that the computer currently has.
- Step 4. Double-click the computer type that is listed under the Computer branch, and then select the **Drivers** tab, **Update Driver,** and then click **Next.**
- Step 5. Click **Display a List of Known Drivers for This Device**, and then click **Show All Hardware of This Device Class.**
- Step 6. Click the appropriate computer type (a computer type that matches the current type, except for multiple CPUs), click **Next**, and then click **Finish.**

9.4.4 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.

Upgrading ATA Hard Disk Drives
This section describes how to upgrade ATA hard disk drives. An example of an ATA drive is shown in Figure 1. To shorten the discussion, the term ATA refers to the following hard disk drives:


ATA Hard Drive

- Integrated Device Electronics (IDE)/ATA
- Enhanced IDE (EIDE)/ATA with Extensions (ATA-2)
- Ultra ATA

The process is the same for all versions of ATA disk drives.

Upgrades to ATA disks generally fall into two categories. They are adding disk drives and replacing existing disk drives with faster or larger disk drives.

Adding ATA disk drives to an existing ATA disk subsystem is relatively straightforward. ATA disk controllers generally have two channels to which ATA devices, such as disk drives, CD-ROM drives, and so on, can be attached. Each channel consists of a ribbon cable that can be up to 18 in. (46 cm) long, to which a maximum of two disk drives can be attached. One end of the channel is attached to the ATA disk controller, which may actually be built in to the system board. The channel, which is a 40-conductor ribbon cable, usually has two 40-pin connectors attached to it. These 40-pin connectors are used to attach ATA disk drives to the ATA channel.

The ATA channels are usually labeled 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:

- The second disk drive can be attached to the same ribbon cable as the existing disk drive using the second 40-pin connector on the ribbon cable. In this case, one disk drive must be set to the master ATA disk role and the other must be set to the slave ATA disk role. Another method would be the user could set cable select (CSEL) for both disk drives.
- The second disk drive can be attached to the secondary ATA channel using a second ribbon cable for the ATA controller. Set this single drive to either the single drive or master ATA disk role depending on the manufacturer instructions for configuring a single disk drive on an ATA channel.

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.

A system basic input/output system (BIOS) upgrade might be required when upgrading from small ATA disk drives to very large ATA disk drives.. See Section 9.5.2, Upgrading adapter BIOS or firmware, for details on how to upgrade the system BIOS. Older system BIOS might not have the capability to address all the space on very large ATA disk drives. The definition of very large has changed over the years. Various definitions of large have been 504 MB, 1 GB, 2 GB, 4 GB, and 8.4 GB. Many of these have proven to be barriers to ATA that had to be overcome by newer, better, or improved BIOS.

Disk drive speeds are characterized by their rotational speed. Common rotational speeds are 5400 rpm, 7200 rpm, and 10,000 rpm. The faster the rotational speed, the faster the disk can access data. Upgrading from slow disk drives to faster disk drives involves a complete replacement of the disk drives.

Upgrading IDE/ATA/EIDE/ATA-2 Disk Drives to SCSI Disk Drives
There is really no upgrade path from IDE/ATA/EIDE/ATA-2 disk drives to SCSI disk drives. To change from IDE/ATA/EIDE/ATA-2 disk drives to SCSI disk drives, the user must remove all the IDE/AT A/EIDE/ATA-2 disk drives, remove or disable the built-in IDE/EIDE controller, and install an SCSI bus controller and SCSI disk drives. Remember that to boot from an SCSI disk drive, the BIOS on the SCSI bus controller must be enabled, and the SCSI ID of the boot disk must be set to 0.

Upgrading SCSI Hard Disk Drives
Upgrades to SCSI disk drives fall into two categories:

1. Adding SCSI disks to an existing SCSI channel
2. Replacing existing SCSI disks with disk drives that have a faster rotation speed

The SCSI disk drive in the upgrade should match the existing SCSI disk in the following ways:

- SCSI level, which is 1, 2, or 3
- Type, which is normal or wide
- Signaling system, which is SE, LVD, or differential

Adding SCSI Hard Disk Drives
Adding SCSI disk drives to an existing SCSI channel is a rather simple process. See Figure 2. 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 3. 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, shown in Figure 4, might also need to be removed. Remember that the SCSI bus must be terminated at both ends.

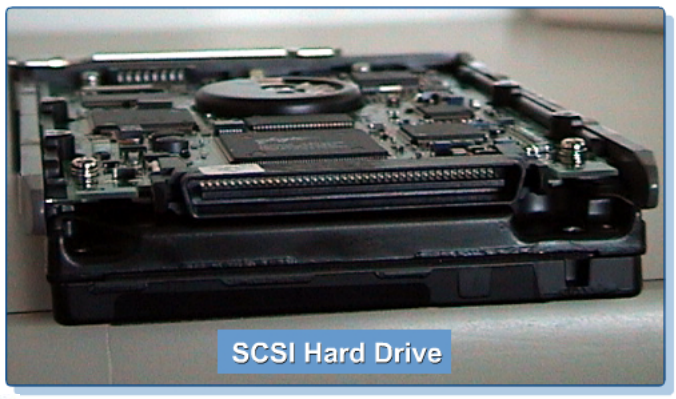

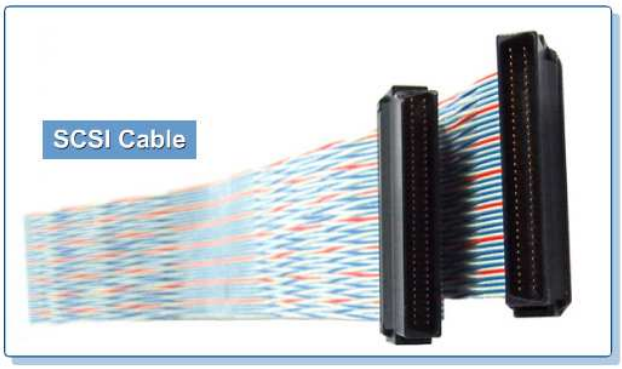

External SCSI devices are usually connected to the SCSI channel in a daisy-chain manner. Adding an additional external SCSI device involves picking an SCSI ID that is not currently in use on the SCSI channel and adding the SCSI device into the daisy chain. Exceeding the SCSI cable length is the biggest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

Replacing SCSI Hard Disk Drives
Replacing an existing SCSI hard disk is just a matter of removing the old SCSI hard disk and checking the SCSI ID. Set the SCSI ID on the new SCSI hard disk to match the SCSI ID of the SCSI disk drive that was removed and install the new SCSI hard disk. If the SCSI disk to be replaced uses an SCA connector, just remove the old SCSI disk drive and insert the new SCSI disk drive in its place. SCA connectors automatically set the SCSI ID of the SCSI disk drive.

Adding 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.

New Drives in a Separate Array
If the newly installed disk drives are configured as a separate array from the existing array, the data on the existing disk

array will be unaffected.

For example, if there is an existing RAID 5 array consisting of three disk drives and there are two new disk drives that need to be installed and configured as a RAID 1 array, the original RAID 5 array will not be affected. No loss of data on the original array will occur. After installing the two new disk drives, use the array configuration utility to initialize the two new disk drives and configure them as a RAID 1 array. The RAID controller will then have two separate RAID arrays configured. They are the original RAID 5 array, consisting of three disk drives and the new RAID 1 array, consisting of two disk drives.

New Drives in an Existing Array
If the newly installed SCSI disk drives need to become part of an existing RAID array, the installer needs to initialize all disk drives in the array, including the existing drives. This means a loss of data on the existing disk drives.

For example, if there is an existing RAID 5 array consisting of three disk drives and there are two new disk drives that need to be installed and configured into a RAID 5 array using all five disk drives, the installer must initialize all five disk drives. The installer can then combine the five disk drives into a single RAID 5 array using all five disk drives. However, the data that was on the original RAID 5 array with the three disk drives will be destroyed in the process and will need to be reloaded from the backup tape. The RAID controller will have a single RAID 5 array consisting of five disk drives.

9.4.5 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.

Check Existing Memory
Before adding memory to a network server, the user needs to verify the current memory configuration. The documentation for the configuration of the network server should have all the details of the memory configuration. However, when this information is not readily available, the user needs to determine it. 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.

Checking Memory Upgrade Feasibility
Before attempting a memory upgrade, first determine whether the network server hardware can support the amount of memory desired in the network server. Some system board control chipsets limit the amount of memory that can be utilized in the network server. Limitations might also apply to the maximum size of a memory module that can be placed into a single memory slot. This will limit the total amount of memory as well. The best source of information about how much memory can be installed in a particular network server is the documentation that came with the network server or the vendor website of the network server.

Verify that the network server operating system supports the amount of memory that is to be installed. Check the vendor website of the network server operating system for maximum memory supported.

Be aware that it might be necessary for the user to remove the existing memory modules and replace them with larger memory modules to achieve the total amount of memory required. An example network server could require four 128-MB modules. This network server may have only four memory slots and two of them could be occupied by 64-MB memory modules. In order to achieve a total of 512 MB, the two 64-MB memory modules must be removed.

Checking Memory Upgrade Compatibility
After answering all these questions, check the vendor hardware compatibility list for the network server to make sure the vendor has certified the memory selected.

One other very important consideration is the metal plating on the leads of the memory module. Two common metals are used on the memory module leads and the connectors in the memory slots. They are tin and gold. Both metals work well. However, never mix the two metals. For example, do not put gold on the memory module and tin in the connector or vice versa. This mismatching of metals causes corrosion at the contact points and results in a bad connection over time. This bad connection causes memory errors to occur.

Installing Additional Memory
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.

Figure 1 shows a diagram of SIMM installation. Figure 2 shows a diagram of DIMM and RIMM installation.

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.

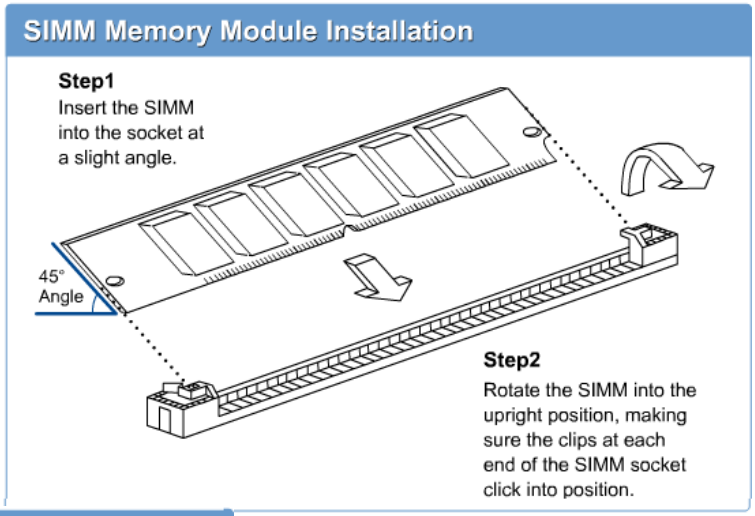

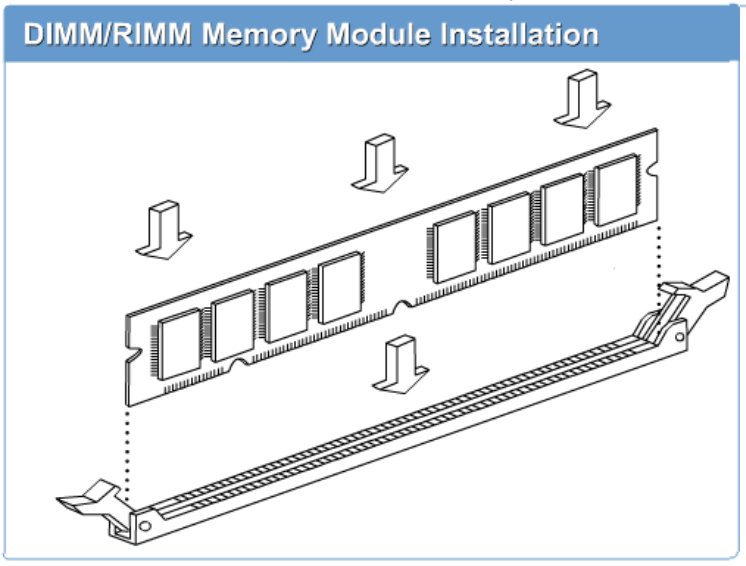

NOTE:
RIMM is commonly believed to stand for RAMbus inline memory module. RAMbus licenses its memory designs to semiconductor companies, which manufacture the chips. Kingston Technology has trademarked RIMM and uses only that term.

While performing a memory upgrade, be sure to follow the upgrade checklist.

9.5 Upgrading Server Components

9.5.1 Upgrading adapter memory

Adapter upgrades also fall into several categories:

- Upgrading components on the adapter, such as memory
- Upgrading the BIOS or firmware on the adapter
- Replacing the adapter with a newer, faster, or more powerful adapter

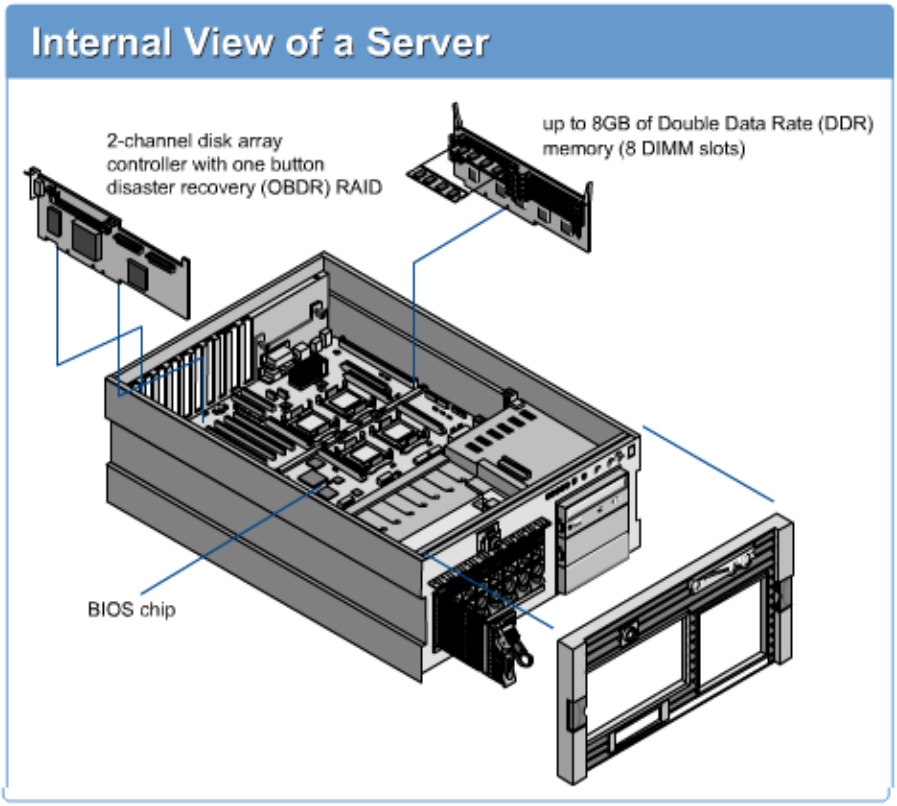

Often, network server adapters have onboard memory that can be upgraded. Follow the adapter manufacturer instructions to perform the memory upgrade. The following adapters use onboard memory:

- **Video adapter onboard memory** – A video adapter uses memory to store the image that is displayed on the monitor. The more memory on the video adapter, the higher the resolution and/or the more colors that can be supported on the video monitor.
- **SCSI adapter onboard memory** – An SCSI adapter uses onboard memory as a buffer or cache between the SCSI disk drives and the network server memory. All read and write operations actually occur to the buffer rather than the disk drive. The larger the onboard memory buffer, the faster information can be supplied to the network server.
- **RAID controller onboard memory** – A RAID controller uses onboard memory in much the same way that an SCSI adapter uses memory as a buffer. One big difference is that the onboard memory on a RAID controller is often backed up by an onboard battery. This prevents the loss of data that is in the buffer if the network server loses power unexpectedly.

9.5.2 Upgrading adapter BIOS or firmware

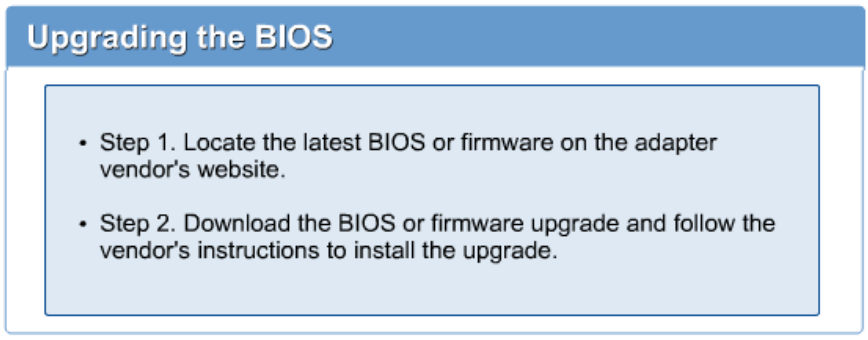

Upgrading the BIOS or firmware on an adapter is very similar to upgrading the system BIOS. The steps are vendor specific, but the general steps are shown in Figure 1.

9.5.3 Replacing an adapter

Adapters are generally replaced after they fail. The replacement procedure generally follows several simple steps, as shown in Figure 1:

Steps for Replacing an Adapter

- Step 1. Power down the network server.
- Step 2. Remove the defective adapter.
- Step 3. Install the new adapter.
- Step 4. Power up the network server.

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.

Steps to Add PCI Hot Adapter

- Step 1. Open the network server chassis.
- Step 2. Open the slot release lever on an available PCI slot. This removes power from the select PCI slot.
- Step 3. Install the adapter into the selected PCI slot.
- Step 4. Attach any necessary cables to the adapter.
- Step 5. Close the slot release lever to secure the adapter in the PCI slot.
- Step 6. Press the PCI hot plug button. This reapplies power to the PCI slot.
- Step 7. The network server operating system locates and loads the appropriate device drivers for the adapter or prompts the installer for the location of the appropriate device drivers.
- Step 8. Close the server chassis.

The steps listed in Figure 2 explain how to use a PCI hot plug to add an adapter.

9.5.4 Upgrading peripheral devices

Internal and External 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. See Figure 1.

Internal peripherals include such components as disk drives, CD-ROM drives, floppy disk drives, and network interface cards. Upgrading disk drives such as EIDE and SCSI is covered in the section 9.4.4, Adding hard drives, in this module. Replacing, upgrading, or adding internal peripherals such as a CD-ROM drive, a DVD-ROM drive, a ZIP drive, a tape drive, or a NIC requires the user to shut down the network server. To install the drives, follow the manufacturer installation instructions. To install the NIC, follow the manufacturer installation instructions after identifying an empty PCI slot into which the NIC will be installed. Be sure to download the latest drivers for the NIC from the website of the NIC vendor. Remember that some

network servers will have peer PCI buses and the "data load" should be balanced among the buses. This requires some knowledge of which PCI slots are on which PCI bus, as well as the data load placed on the PCI bus by the adapters currently on each PCI bus.

If the adapters are not plug-and-play, the user may have to configure the adapters with an interrupt request (IRQ), a direct memory access (DMA) channel, and input/output (I/O) address. The IRQ, DMA, and I/O address cannot conflict with any adapter already installed in the network server. As with all upgrades, be sure to follow the upgrade checklist.

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. Follow the manufacturer installation instructions for external peripherals. As with all upgrades, be sure to follow the upgrade checklist.

9.5.5 Upgrading system monitoring agents

System monitoring agents are software specific to the network server itself and supplied by the network vendor. These agents are generally installed at the installer's option when using the vendor-supplied operating system installation-assistance software. As the network server vendor releases newer versions of the system monitoring agents, the system monitoring agents must be upgraded. Because the system monitoring agents are extremely vendor specific, the user must follow the installation instructions of the vendor to upgrade the system-monitoring agents. In some cases, the network-monitoring agents can predict the impending failure of network server components such as the processor, the memory, or the hard disk drives.

System monitoring agents monitor various aspects of the network server such as configuration, mass storage, network interface card, system utilization, thermal conditions, and operating system status. The agents usually use standard protocols such as hypertext transfer protocol (HTTP), simple network management protocol (SNMP), and desktop management interface (DMI) to report information to a management console. The management console might be a standard Web browser such as Netscape or Internet Explorer, a vendor-supplied management console such Compaq Insight Manager, or a third-party network management console such as Hewlett-Packard OpenView. Figure 1 shows the box for HP OpenView.

As with all upgrades, follow the upgrade checklist.

9.5.6 Upgrading service tools

Due to the rather unique nature of network servers compared to standard desktop microcomputers, a wide variety of service tools are sometimes necessary to configure, troubleshoot, and maintain them. These service tools are software that is installed on the network server during the installation and configuration processes.

Network server vendors often release new versions of these tools to fix bugs, add new features, or add support for new hardware to the utilities.

Some service tools are part of the network operating system. Updates to these service tools are available from the network server operating system vendor.

These updated tools are generally available on either CD-ROMs supplied by the network server vendor or from the vendor website. After having access to the updated tools, follow the installation instructions of the vendor to install the updated tools on the network server. As with any upgrade, follow the upgrade checklist.

Upgrading service tools is generally a software upgrade operation, and, the upgrade checklist should be followed.

Upgrading the utilities on the diagnostic partition, which is only available at network server boot time, usually requires the user to shut down the network server. This means that the user needs to schedule downtime to upgrade the software utilities on the diagnostic partition.

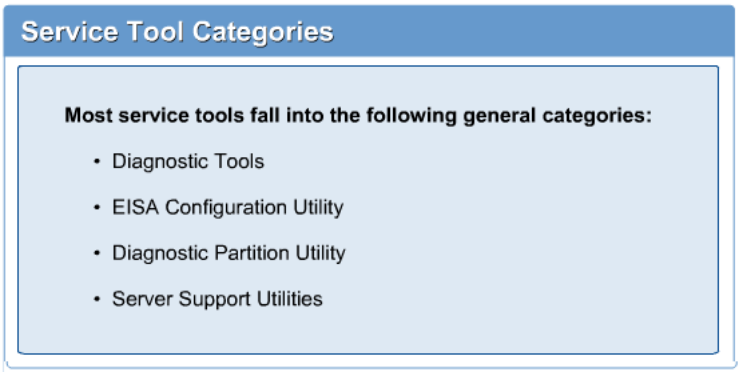

Most service tools fall into the following general categories:

- **Diagnostic Tools** – Several sets of utilities can be categorized diagnostic tools. Some diagnostic tools might be installed on the diagnostic partition and specific to the network server. Other diagnostic tools comprise part of the network server operating system. There are also third-party diagnostic tools.
- **Extended Industry Standard Architecture (EISA) Configuration Utility** – On network servers that have an EISA bus, the EISA configuration utility enables the user to configure the components of the network server. EISA is not plug-and-play, and adapters installed into the network server must have a configuration file loaded from a disk to allow configuration of the adapter. The EISA configuration utility is used while the network operating system is not loaded.
- **Diagnostic Partition Utility** – The utilities in the diagnostic partition enable the user to view and change the configuration of the hardware components in the network server without the network server operating system being loaded. This enables the user to troubleshoot nonfunctioning hardware components.
- **Server Support Utilities** – These utilities such as backup software and antivirus software play a vital role in support of the network server.

9.5.7 Document the configuration

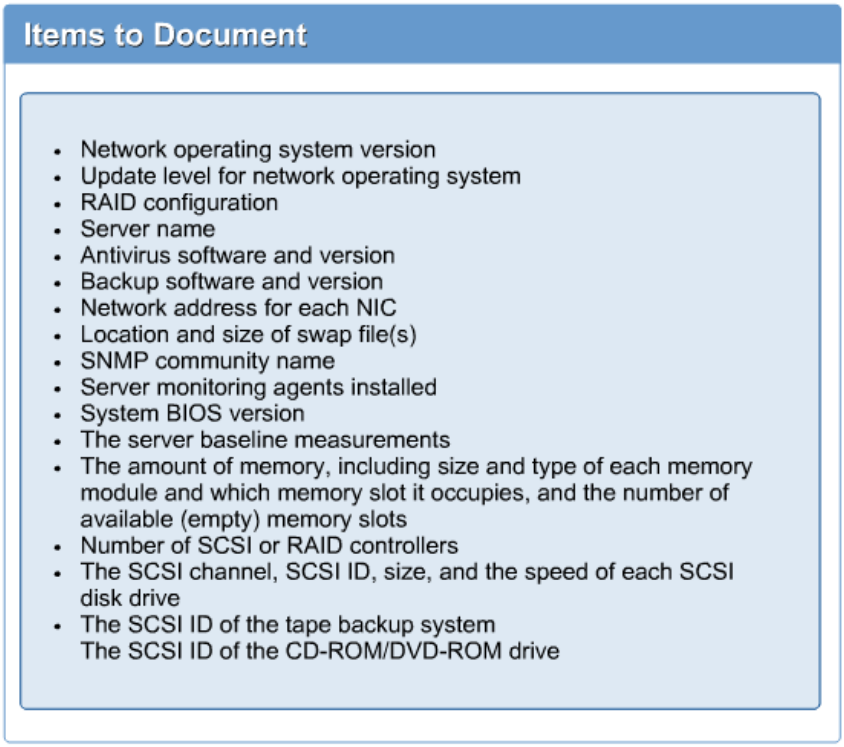

A very important part of being a good administrator is to document the problems and procedures to fix them. The reason good documentation is needed is so that if a problem occurs more than once, the user can look up how to fix it and prevent having the same failed method twice. This will in turn save the time fixing it, and the user will be able to have the server back online much faster.

NOTE:
Documentation is a part of the upgrading and troubleshooting processing. Be able to explain the documentation process and why it is so important.

As the user installs, configures, or fixes the server, the user should record exactly step by step the success and failure of what is done. The server also has logs of its own. These logs are stored on the hard disk and can become rather large. Depending on the amount of hard disk space delegated to the log files, the user may have to print or archive them to backup media. These files need to be removed to make room for the logs yet to be created by the

server. By doing this documentation process, the user now can quickly troubleshoot the server when a problem arises.

Figure 1 lists the items to document during and after the server is installed and configured.

Summary

This module discussed Advanced Hardware Fundamentals for Servers. Some of the important concepts to retain from this module include the following:

- Fault tolerance is the ability for a system to continue when a hardware failure occurs. Fault tolerance is important in a network server that provides users the ability to share files, programs, printers, and so on.
- Redundant Array of Inexpensive Disks (RAID) is designed to provide fault tolerance in the event of a disk drive failure on a network server.
- RAID 1 requires two disk drives. All other RAID levels except RAID 0 require at least three disk drives. There are two ways to implement RAID 1. That is through disk mirroring or through disk duplexing.
- RAID is typically implemented using a RAID disk controller but can also be implemented in software. Software RAID is implemented at the disk partition level supporting RAID 0, 1, and 5. Hardware RAID is implemented on the physical disk providing more reliable fault tolerance.
- External disk subsystems are added when the amount of disk storage cannot be accommodated by internal disk drives. Simple external disk subsystems with only a few disk drives function in the same way that internal disk drives function. Large subsystems can have their own RAID mechanism built in and are often configured separately from the network server.
- Upgrading the processor or installing an additional processor can improve the performance of the server. Before proceeding, verify the motherboard can accommodate the upgrade or supports multiple processors. Check the motherboard manual or the website for the motherboard manufacturer.
- Disk drive upgrades can be accomplished by adding to the existing drives or replacing drives with larger and faster drives.
- Increasing memory on the server will improve performance but processors have a maximum that they can support. Check the feasibility and compatibility of adding new memory to the existing server before proceeding.
- System monitoring agents and service tools should be updated regularly. System monitoring agents monitor configuration, mass storage, the NIC, system utilization, thermal conditions, and the operating system status. Service tools are used to maintain the system and for troubleshooting purposes.
- Documenting the configuration of the network server is important to a network administrator. Service logs and other log files provide valuable data needed to troubleshoot the system if problems occur. Make sure information is detailed and current.

The next module introduces the student to PC Networking. It discusses the types of networks, the components of a network, and connecting to the Internet.

Quiz

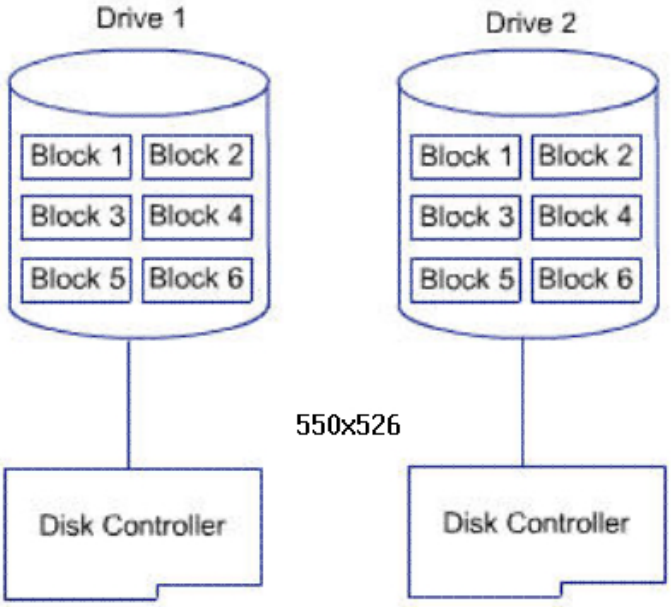

Which term best describes the figure?

- disk mirroring
- striping, no parity
- striping, with parity
- ✓ disk duplexing

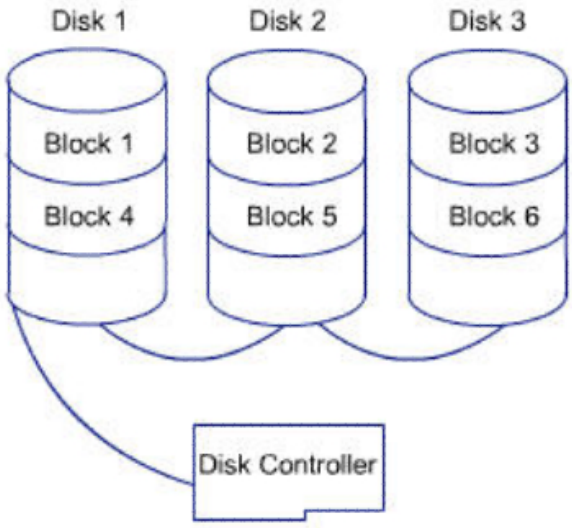

Is the RAID array shown in the graphic configured for fault tolerance?

- Yes, blocks 1 and 2 have identical data, blocks 3 and 4 have identical data, and blocks 5 and 6 have identical data.
- Yes, blocks 1 and 4 have identical data, blocks 2 and 5 have identical data, and blocks 3 and 6 have identical data.
- ✓ No, each block shown has different data.
- No, the array requires five disks to accommodate fault tolerance.

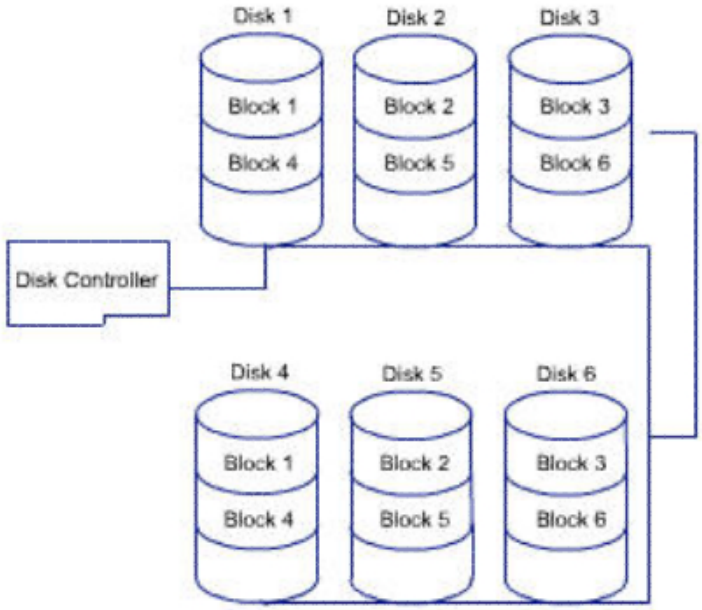

Which term best describes the figure?

- RAID 0
- RAID 1
- RAID 5
- ✓ RAID 0/1

Correctly order the seven steps required for installing a new processor.

- ✓ Follow the upgrade checklist.
- ✓ Upgrade the system BIOS.
- ✓ Open the network server chassis using ESD best practices.
- ✓ Remove the current processor.
- ✓ Insert the new processor.
- ✓ Close the network server chassis.
- ✓ Verify the new processor is recognized by the network server hardware and the network operating system.

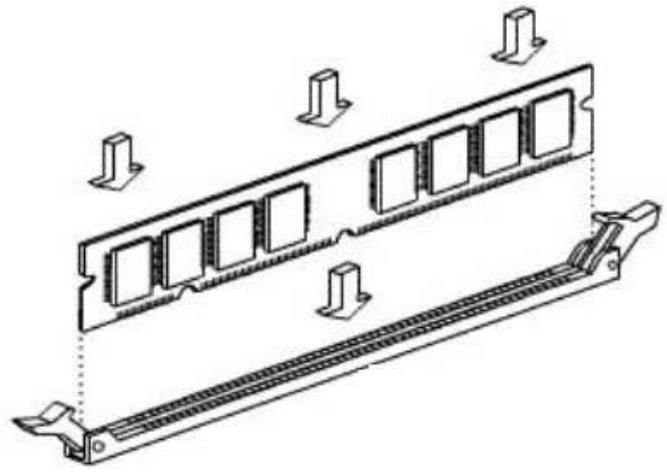

Which type of RAM does the procedure indicate? (Choose two.)

- ✓ [x] DIMM
- [ ] SIMM
- [ ] SRAM
- ✓ [x] RIMM

With a RAID 5 array of three disks, how many disks can fail without the data in the array being permanently lost?

- ( ) 0
- ✓ (•) 1
- ( ) 2
- ( ) 3

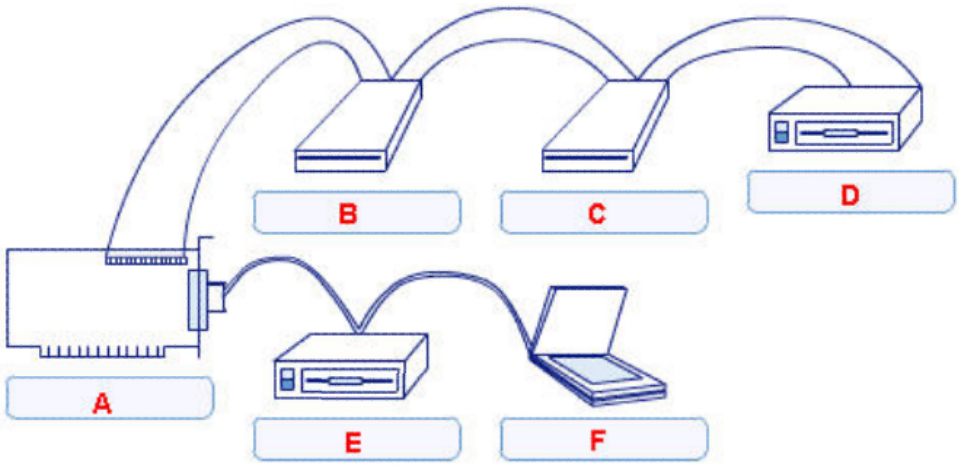

Which locations on the SCSI chain should have termination? (Choose two.)

- [ ] A
- [ ] B
- [ ] C
- ✓ [x] D
- [ ] E
- ✓ [x] F

What is the total storage capacity of a RAID 5 array that has five 30 GB hard drives?

- 30 GB
- 90 GB
- ✓ 120 GB
- 150 GB

What is the maximum number of disks that can be attached to an ATA RAID controller?

- 1
- 2
- ✓ 4
- 5

Which methods can be used to connect an external disk subsystem to a network server? (Choose two.)

- ATA controller
- ✓ SCSI cable
- USB port
- ✓ Fibre Channel