Server+ Server Hardware Specialist

Version 3.1.0

Server+ (Server Hardware Specialist)

Version 3.1.0

Abstract:

This study guide will help you prepare for the CompTIA exam, Server+. Exam topics include Advanced Hardware Issues, such as RAID, SCSI, Multiple CPUs, SANs, Server Types, System Bus Architectures, and Disaster Recovery, Upgrading, and Security Concepts.

Find even more help here:

Contents:

SERVER+ (Server Hardware Specialist)

Who is a Server Hardware Specialist?

(From CompTIA)

- Has in-depth knowledge of servers, including working knowledge of troubleshooting, physical security and disaster recovery. Can recover from a server failure.
- Ensures high availability by meeting the Service Level Agreement requirements, including proactively recognizing and responding to problems and performing recovery.
- Has thorough working knowledge of hardware configuration and network connectivity. Also has the ability to perform problem determination for all aspects of the server (hardware, software, networking).
- Installs and configures server hardware to meet application requirements
- Implements current and emerging data storage and transfer technologies such as SCSI, RAID, FC-AL.
- Has a thorough working knowledge of networking protocols (e.g., TCP/IP, IPX/SPX, SNMP, DMI, SNA) for diagnosing the impact of the network on the server and vice versa.
- Provides support, including second-level support, for resellers and end-users.
- Performs maintenance on server systems, data-storage subsystems, and network devices.
- Has good planning and integration skills and is able to:
    - Upgrade a server without impacting network users
    - Increase storage capacity without impacting network users
    - Design and implement a data recovery plan in the event of a network device failure
    - Perform peripheral upgrades, BIOS upgrades, memory upgrades, processor upgrades, mass storage upgrades, and adapter upgrades
- Demonstrates high levels of leadership through mentoring and training others in server concepts and operations.

Definition and types of Servers

Server – Generally, it's a computer on a network that manages resources on that network. Servers usually are identified by what they do. For instance, a network server manages network traffic. A file server stores user files. An application server stores applications for users to use

- A server can be **general purpose**, which means it can execute several programs at once and can be a client as well as a server. In this case, "server" refers to a program used to manage resources (such as a database), not the entire computer.
- A **dedicated** server performs no other tasks but server tasks. In this case, "server" refers to the entire computer, not just a program. For instance, Microsoft's database and mail servers are SQL and Exchange server. They run as separate applications on the Windows NT or Windows 2000 Server Operating System platforms.

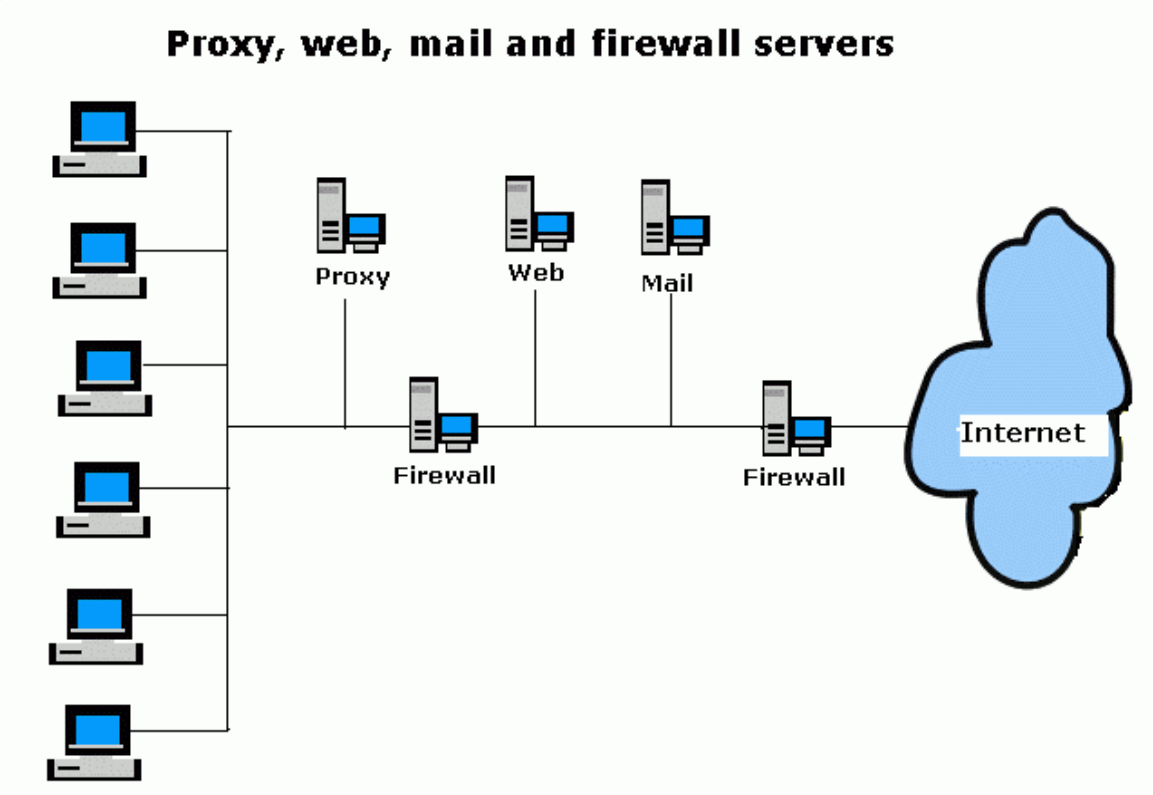

SERVER TYPE	DETAILS
Gateway Sun White Paper	Translates from one network communication type (i.e., protocol) to another. An example would be Microsoft's SNA (see below). A server can become a gateway when you add another NIC to make it multihomed.
Router	This is a server configured to operate as a router. It is not a hardware router. Moves packets between network segments at Layer 3 of the OSI Model. Like a gateway server, a router server requires at least two NICs to make it multihomed. The option also must be enabled in the server's OS. Compare to the server that acts as a bridge.
Bridge Bridge basics	Joins network segments and manages packet traffic between segments. Works with packets at the Data Link Layer (Layer 2) of the OSI model by using the MAC address hard coded on the NIC. Good for connecting dissimilar LANs, like Ethernet and Token Ring.
Firewall 3COM Tutorial Cramsession article on ISA 2000	Prevents unauthorized access to or from a private network. This is a software-based firewall, in contrast to the base IOS a router uses. This server should be at least dual homed but it can be Tri-homed (three NICs). This is a way to create an internal network, a DMZ and an external network.
Proxy **Microsoft products** Proxy Server Proxy direct access **Novell products** BorderManager	Acts on behalf of the client's request to access a web server. Uses filters to permit or deny access to certain sites. Proxy will cache frequently visited pages, which allows the page to be accessed faster upon revisit. Also can segregate segments within a private network, thus creating a DMZ.
Database SQL Server	A database server has a distributed database. It stores a database for many clients to access, and generally has **VERY** high

SQL Direct Access	hardware requirements, especially RAM and hard drive. An example would be an SQL server or SQL Direct Access server. Databases can be relational as well as distributed. One server can serve more than one single database. Databases are generally accessed by applications, not directly by clients

Client For a Client/Server FAQ, click here. This page is not being actively maintained but the contents are excellent.	Client/Server is a type of architecture. It is more prevalent than database servers. Clients access resources available on the network thru this server. An FTP server (see below) is a typical example.
Application TechMetrix article	An application server is designed to make it easier for developers to develop three-tier applications. This type of server stores applications for clients to use and simplifies application management because software is installed only once. Three types: peer-to-peer, distributed and dedicated.
Mail Microsoft Exchange server Novell's GroupWise	Handles email function. When connected to the Internet, it will use Simple Mail Transfer Protocol (SMTP) /Multi-purpose Internet Mail Extensions (MIME).
FTP	File Transfer Protocol. Allows the movement of one or more files between computers. Generally includes security and data integrity controls appropriate for the internet.
SNA Intro to SNA	System Network Architecture. Two words: IBM and mainframe. This server acts as a gateway between a PC/server-based network and a mainframe.
RAS	Remote Access Server. Allows remote users access to the LAN. Good acronyms to remember: POTS (Plain Old Telephone Service) and PSTN (Public Switched Telephone Network). There is no difference between the two.
File and Print	Oldest and most common server type. Provides clients access to file and print resources.

Fax Fax for Windows NT/2000	Allow the transmission, receipt and delivery of faxes. This type usually is integrated with an email server.
DNS	Domain Name Service/System. Maintains a database of records that allow, among other things, the mapping of IP addresses to domain names. Configured correctly, a DNS server will resolve fully qualified DNS names (www.cramsession.com) to an IP address (206.112.74.39).
WINS	Windows Internet Name Service. Configured to resolve NetBIOS names to logical address. This type of server is required in an NT and W2K mixed mode environment.
DHCP DHCP overview	Dynamic Host Configuration Protocol. Clients boot and seek out this server to receive an IP address and network access. The address can expire because of the lease. The DHCP server itself must have a static IP address, which means the server itself cannot be a DHCP client.
WEB	A Web Server delivers client requested web pages, usually in conjunction with a Proxy server. Can refer to the server itself or the software that carries out the service. Examples include Apache and IIS. For a white paper on web server basics, click here.

For server types, click here.
For vendor-specific products: Novell | Microsoft | Unix | Linux | IBM | Oracle.
Cramsession on CIW Server Administration, click here.

Physical vs. Logical drives

A **Physical** drive is one you can pull out of the box. If you look inside My Computer, the first physical drive will be assigned the letter "C". Add another physical drive and it'll get the next drive letter. In theory, you can have up to drive Z. A physical drive can have more than one logical drive.
A **Logical** drive is a partition on a physical drive. A physical drive can be segmented into multiple logical drives. When the corresponding physical drive comes out of the computer, so does the logical drive.
Remember . . .

- Always place the swap file on a physical drive with plenty of space. Placing the swap file on a logical drive can degrade server performance.
- As a rule, if two or more physical drives are present, the physical drive will always take precedence in drive letter assignments. If you have two physical drives and you create a logical drive on the first physical drive (C:), it would still favor the second physical drive for D: and the logical drive would receive E.
- In Windows2000, use the disk management MMC to change drive letters to anything from C-Z.

For a good explanation of how to Use the Fdisk Tool and the Format Tool to Partition or repartition a Hard Disk, go here:
http://support.microsoft.com/support/kb/articles/Q255/8/67.ASP

SCSI vs. EIDE/ATA hard drives

EIDE/ATA (Enhanced Integrated Drive Electronics/Advanced Technology Attachment) is generally less expensive and more readily available than SCSI (Small Computer System Interface) drives. SCSI drives also are slightly more complicated to install than IDE/ATA, because of termination requirements. However, EIDE/ATA is limited in its lack of redundancy (lose one device, the entire system is destroyed: see the RAID section later) and in the number of devices that can be controlled from a single paddle card. EIDE/ATA is limited to hard drives and not many at that. SCSI, in addition to hard drives, can handle CD-ROMs, scanners, and tape drives, either inside or outside the case (EIDE/ATA is strictly inside the computer, and supports no external peripherals at all).

For EIDE/ATA versus SCSI white paper by Dell, click here.

For an overall breakdown of EIDE/ATA & SCSI hard drives, click here.

Caution: do not confuse ATA (Advanced Technology Attachment) with ATAPI (ATA Packet Interface). ATAPI is the IDE interface that works between the computer and either CD-ROM devices or tape backup drives.

SCSI Overview

- SCSI – SMALL Computer Systems Interface.
- Since SCSI uses Bus architecture, multiple devices can share it.
- Host Bus Adapter (HBA) is the card that connects your System Board to the SCSI bus.
- The SCSI Host Bus Adapter must be compatible with the system's bus.
- Disabling the motherboard BIOS IDE is common if you need to boot off the SCSI device running on a chain off the Host Bus Adapter.
- SCSI Systems can be both internally and externally mounted.
- SCSI has a disconnect/reconnect/reselect feature that allows a device to temporarily give up control of the SCSI bus.
- Both ends of the SCSI bus require termination.

SCSI basics . . .

SCSI Technology	Details
SCSI-1 - The "Original" SCSI specification released in 1980. It generally was compatible with difference SCSI implementations. - Narrow and regular.	8-bit parallel interface between host adapter and device. Runs at 5 MHz and has a max data transfer rate of 5 MBps. Does not support SCSI parity.
SCSI-2 - The beginning of the SCSI standard. Improved the data transfer rate and supported new devices.	Enhancements include newer command sets, wider data paths (16 to 32 bits) and command queuing, scripting and disconnects. 5 to 10 MBps data transfer rate.
SCSI-2 FAST	"Fast" SCSI-2 doubles the data transfer rate over the existing data path by doubling the clock rate from 5MHz to 10MHz. "Fast" SCSI works on either single ended or differential (see below). 8-bit bus with a data transfer rate of 10 MBps.
SCSI-2 WIDE	Doubles the original 8-bit bus width to 16 bits. Doubles peak transfer rates to 10 MBps. Uses an extra cable or (more commonly) a 68-pin P cable. Up to 16 devices per bus.
Ultra SCSI (AKA, FAST20)	Increases bus speed from 5 MHz to 20MHz. 8-bit bus with a 20MBps transfer rate.
Ultra Wide SCSI	Doubles the bus width of Ultra SCSI, which doubles the peak transfer rate. 16-bit bus with a 40MBps Transfer rate.

Ultra-2 SCSI (AKA, FAST40)	Increases bus clock to 40MHz but **requires** differential signaling and LVD bus drivers (see below).
Wide Ultra-2 SCSI	Doubles the bus width of Ultra-2 SCSI to double the transfer rate to achieve 80 MBps.
SCSI-3 (AKA Ultra SCSI or Ultra-3 or Ultra160)	Doubles the SCSI clock speed to 80MHz, which doubles the transfer rate to 160 MBps. **Requires** LVD bus drivers.

SCSI Technology	Max Cable length (in meters)	Max Speed (MBps)	Max number of devices
SCSI-1	6	5	8
SCSI-2	6	5 to 10	8 or 16
FAST SCSI-2	3	5 to 10	8
WIDE SCSI-2	3	20	16
FAST WIDE SCSI-2	3	20	16
Ultra SCSI-3, 8-bit	1.5	20	8
Ultra SCSI-3, 16-bit	1.5	40	16
Ultra-2 SCSI	12	40	8
WIDE Ultra-2 SCSI	12	80	16
Ultra-3 (Ultra160)	12	160	16

Don't forget: 0-7 = 8 devices and 0-15 = 16 devices.

SCSI Installation Steps:

1. Choose the Host bus adapter (HBA) / SCSI system.
2. Install the HBA and SCSI devices into the server
3. Assign SCSI IDs to the peripherals (set jumpers)
4. Enable or disable SCSI parity on the peripherals.
5. Properly place and install your SCSI cabling.
6. Terminate BOTH ends of the SCSI chain.
7. Test the configuration.

*** What you won't have to worry about ***

1. Master and slave relationships
2. The number of heads on cylinders.

Though less common today, computers with older system BIOS will not recognize an IDE hard drive until the number of heads and cylinders are keyed in.

*** What to remember ***

- Add one device at a time and test.
- Use good termination. Use the best available terminator.
- Document all settings.
- Use high-quality cables. The cables should be rated for SCSI and should be the same make.
- Don't use longer cable lengths than the bus allows.

SCSI Bus Width (Narrow and Wide) and Bus Lengths:

- Wide SCSI is 16-bit and has a greater bus width that produces greater bus throughput.
- Narrow SCSI width is 8-bit and has less throughput.
- Adapters must be used when mixing narrow and wide devices together in an external daisy chain.
- Remember this rule: Wide goes first.
- Use all your wide devices first and closest to the Host Bus Adapter.
- After placing the last wide device, use adapters to go from 68-pin to 50-pin connections to accommodate narrow devices.

SCSI External Cables

External SCSI devices Cable	The device will have a 50-pin female Centronics connector. The cable will have a 50-pin male connector.
SCSI peripheral cable (50-pin Centronics) Two types: 1) Thick and 2) Thin	50-pin Centronics connectors on both ends of the cable. Is used to connect between SCSI-1 host adapters and SCSI devices or to "Daisy Chain" from SCSI device to device.
DB-25 to Centronics 50 (SCSI-1 Cable)	Older SCSI-1 cables use a male DB-25 connect on one end and a Centronics 50-conductor connector on the other end.
HP 50 Connector	With newer SCSI adapters, do not use the Centronics connector for external devices. Instead, use a miniature 50-socket D shell connector.
Mini DB 50 to Centronics 50 (SCSI-2 Cable)	A common cable with a mini DB50 on one end and a Centronics 50 conductor connection on the other end.

SCSI Internal Cables

- Use a single, 50-in head-connector connection.
- Internal cables are ribbons cables with 50-pin IDC connectors (see below)
- The connectors sit in the middle of the cable.

Connectors

50-pin (Internal)	Two types: 50-pin NON-shielded connector alternative 2 (recognized a typical IDE cable with 50 pins). 50-pin NON-shielded connector alternative 1 (recognized as a connector with a d-shaped male receptacle that protects its pins).
50-pin (External)	50 pins arranged in two rows of 25 and resemble a Centronics printer connector. It's a shielded connector and is the most common of the external connector types. Other types resemble a D shaped female connector with 50 pin slots. Both a shielded and an unshielded version.
68-Pin	Wide SCSI will require more connections than are available with 50-pin variations. 68 pins will accommodate Wide SCSI. Both a shielded and an unshielded version.
80-pin SCS (Single Connector Attachment)	Used for hot swappable drives and removable drives that slide into bays.

25-pin variations were common with Apple Macintosh computers:

- The single-ended 50-pin cable is reduced to 25 pins by tying grounds together.
- Uses a DB25 connector.
- Often used as an external SCSI connector.

SCSI IDs

- A DIP switch or jumper usually sets IDs on devices.
- IDs must be unique on the chain, no two can be alike.
- Devices are distinguished from each other on a bus by SCSI IDs.
- On SCSI the numbered ID ranges are 0-7 on SCSI-1 or SCSI-2, 0-15 on a single Channel SCSI-3 and 0-31 on a Dual Channel SCSI-3.
- Priority is highest to lowest, 7-0 / 15-8.
- Usually, the slowest device should be set to the highest number of all the devices other than the controller. Set the controller to the highest ID.

SCSI LUNs (Logical Unit Number)

- Each SCSI device can have sub-devices, or LUNs.
- In SCSI-2 specifications, each SCSI device address can have up to eight LUNs, numbered 0-7.
- A LUN can, itself, be divided into sub devices. Called LSUNs (Logical Sub Unit Number), these can range from 0 to 255 devices.
- An example of LUNs in action is an optical disk jukebox.
- Some HBAs might ignore LUNs unless the "Enable LUNs" option is set in the host adapter BIOS.

SCSI Termination

A SCSI bus must be terminated in two – count 'em: two and only two – positions -- at the beginning and at the end.

Three types of termination

Passive	Pairs of resistors. A 220 Ohm pulls the signal up to TERMPWR and a 330 Ohm pulls the signal down to GROUND. Simpler and cheaper than Active or FTP. Best with four or fewer devices. Adequate for SCSI-1. In SCSI-2, they're called "Alternative 1" but active terminators are recommended wherever possible for Single Ended SCSI (see below). Only good on distances of 2-3 feet. Passive termination is ALWAYS used with differential (HVD) SCSI (see below).
Active	Voltage regulator ensures the correct termination voltage. 110 Ohm resistors are connected from each signal line to a common 2.85 Volt regulated power supply. Terminates the bus better than Passive and supplies cleaner pull. First used in SCSI-2, in which they're called "Alternative 2." Many have an LED showing the termination level.
Forced Perfect Termination (FPT)	Single Ended termination method. Uses diode switching and biasing to make up for any impedance mismatch between the SCSI cable and peripheral device FPT works only with FPT.

SCSI Signaling Systems

Single-ended (SE)	Most SCSI implementations are single-ended because they are cheaper than HVD or LVD. Cheaper wire, fewer pins and simpler electronics than HVD. Most SCSI systems use a single-ended or "unbalanced" electronic signaling system. Uses open collector drivers to drive the SCSI bus. Can handle an improper cable insertion and can detect connection of HVD devices, thus preventing major damage. More prone to picking up interference and noise, so the cable length limit is 6 meters (about 20 feet). Shortest allowable cable length between devices is about 12 inches.
Differential or High Voltage Differential (HVD)	Not compatible with single-ended devices. More expensive than SE. Uses "balanced" or "differential" signals, so twice as many wires are required in the cable. Uses twisted-pair wiring. Uses two wires to drive a single signal. Uses its own return line that is isolated from the reference ground. Used in less than 5 percent of all SCSI implementations. Supports throughput of 40MBps at cable lengths of up to 25 meters (about 82 feet). For this reason, HVD is generally used where long cable runs are necessary, especially in high noise environments.

Low Voltage Differential	Low Voltage Differential Signaling is a low noise, low power for high-speed data transmission. LVDS uses a dual wire system, and thus filters noise more easily and effectively. All Ultra2 SCSI devices use LVD signaling with performance of: - Speed: 40 MBps (8-bit) or 80 MBps (16-bit) - Cable length: up to 12 meters. - Number of devices: 16 max Used to solve incompatibility issues with SE and HVD devices because HVD and SE devices **CANNOT** co-exist on a bus. Can switch to "SE mode" and operate with SE devices on the same bus. Easily integrated with current hardware and firmware and is fully backward compatible. Integration performs faster driver buffer-to-host transfers. Benefits of LVD: - Better reliability - Backward compatibility through automatic mode sensing of single-ended (SE) or differential bus-configuration types. - Common-mode noise immunity. - Reduced EMI.

Remember:

- SE and HVD are incompatible.
- Adapters do exist that will "mate" SE and HVD but it's not considered a best practice. Doing so can damage the HBA or the device itself.
- The SCSI connector does not distinguish between signal types used to make the connection.

See Cramsession's guide on SCSI Part I, Part II and Part III.
Here are very good links to SCSI connectors, complete with pictures.

RAID

- A Redundant Array of Inexpensive (or Independent) Disks -- The name comes from a time when "inexpensive" was, indeed, a relative term. Inexpensive was, originally, the correct term. However, over the years, many techs came to accept Independent as well. Today, both terms are acceptable.
- Data storage system in which data and information needed for error correction are stored over two or more hard drives. RAID provides improved reliability and performance.
- Two types of RAID:
  - 1) Hardware - with a special dual ported SCSI adapter. More expensive than software RAID but provides better performance.
  - 2) Software- integrated into the OS (usually RAID 0 & RAID 1).
- RAID offers
  - High Availability – ensures server resources are available a high percentage of the time.
  - Fault Tolerance - if hard drive in the RAID array fails, another will take its place and data can be restored from other disks in the array. The obvious exception to this rule is RAID 0, which offers no fault tolerance at all.
- A RAID array appears to the OS as a single, logical hard disk.
- Most RAID levels use the technique of striping. Striping is the partitioning of each drive's space into units that can range from a sector (512 bytes up to many megabytes).
- The stripes of the disks will be interleaved and addressed in a specific order.

The minimum number of drives required for each RAID level:

RAID 0	2 drives (but is not redundant)
RAID 1	2 drives
RAID 1+0 (AKA, RAID 10)	4 drives
RAID 2	NA
RAID 3	3 drives
RAID 4	3 drives
RAID 5	3 drives
RAID 6	3 drives
RAID 7	NA
RAID 53	5 drives

RED = More likely to be tested on these RAID levels.

RAID 0	Cheapest and fastest of all RAID technologies. Disk striping without parity. Offers best performance but no fault-tolerance, which means it's not really RAID. Lose just one drive and all the data in the array is lost. (In this case, reinstall and restore data from a backup.) *Never use this in mission critical environments.* Most often seen in video production and editing, image editing, Pre-Press applications and any application that requires high bandwidth. Also useful for proxy caches where fault tolerance isn't so much of an issue.
RAID 1	Disk Mirroring (one controller; if it fails, both hard drives will fail) or Duplexing (two controllers, so more reliable than mirroring). No striping. What is written to the first hard drive will be written to the other. So if one drive fails, replace it and copy data from the other drive. If both drives fail, they both must be replaced and data restored from a backup. Read performance is improved since either disk can be read at the same time. Cost per MB is high. If you mirror/duplex two 30-GB hard drives, you effectively have only 30 GB total. May not support hot swap when implemented through software.
RAID 1+0 (AKA RAID 10)	A mirrored array with segments that are RAID 0 arrays but offers the same fault tolerance as RAID 1. Same fault tolerance overhead as mirroring. Offers higher performance than RAID 1 but at a much higher cost. Excellent choice in environments where high performance is required but achieving maximum reliability is not a concern. If one drive fails, the entire array becomes, essentially, a RAID 0 array.

RAID 2	Striping across disks with some disks storing Error Checking and Correcting information. Called "On the fly" data error correction. Slow. Very little fault tolerance. Relatively simple controller design compared to RAID levels 3,4 & 5. Installation costs can be very high. There are no commercial implementations; not commercially viable.
RAID 3	Offers striping. One drive is dedicated to storing parity. Cannot overlap I/O and, therefore, is best for single-user systems with long record applications. Very high Read/Write data transfer rates. Controller design is relatively complex. Very difficult and resource intensive to implement as software RAID.
RAID 4	RAID 4 is similar to RAID 3. The difference is an entire block of data is written to each drive, rather than individual bits, as in RAID 2 and RAID 3. Uses large stripes, so records can be read from any single drive. All write operations must update the parity drive and no I/O overlapping is possible.

RAID 5	Striping with parity. Data and parity information is striped at the block level across all the drivers in the array. Most commonly used of all the RAIDs. Commonly implemented in enterprise critical file and application servers; data base servers; Web, email and news servers and intranet servers. Fastest and most reliable of RAID technologies. Requires a minimum of three hard drives but usually five or mores disks are used. If one drive fails, parity information stored on other drives can be used to reconstruct data.
RAID 6	An extension on RAID 5, it includes a second parity algorithm in another stripe. Parity scheme requires N+2 drives. More fault tolerant than RAID 5 but it is also more expensive.
RAID 7	Multiple striped disks each with its own controller. Each stripe has its own parity drive, which offers very fast performance Very expensive. Access times decrease with each increase in the number of actuators in the array No extra data transfers required for parity manipulation. Only one vendor currently is offering this RAID level technology (Storage Computer Corporation. For more information, click here and here.)
RAID 53	Can be thought of as "RAID 03" because it is implemented as a striped array (like RAID 0) RAID 3 array segments. RAID 3 array segments offer high data transfer rates. RAID 0 striping offers high I/O rates for small requests.

RED = Most likely to be tested on these RAID levels.
For an excellent Cramsession article on RAID technologies, click here.

EIDE/ATA (Enhanced Integrated Drive Electronics/Advanced Technology Attachment)

You might see it as AT, IDE, EIDE, or spelled out as it is here. Though there are obvious technical differences, all the terms generally refer to the same thing, a standard electronic interface used between a computer motherboard's bus and the computer's disk devices.

EIDE/ATA vs. SCSI

As previously stated, EIDE/ATA (Enhanced Integrated Drive Electronics/Advanced Technology Attachment) is generally less expensive and more readily available than SCSI (Small Computer System Interface) drives. EIDE/ATA devices do not have the same termination requirements as SCSI drives, so are not as complicated to install. However, EIDE/ATA is limited in its lack of redundancy (lose one device, the entire system is destroyed) and in the number of devices that can be controlled from a single paddle card. EIDE/ATA is limited to CD ROM drives and additional hard drives and not many of either. SCSI, in addition to hard drives, can handle CD-ROMs, scanners, tape drives, either inside or outside the case (EIDE/ATA is strictly inside the computer, and thus supports no external peripherals at all).

EIDE/ATA overview

- Supported directly by system BIOS in most cases.
- Extends the data-transfer rate up to 6.67 MBps, or 10 MBps for a cache-hit burst.
- The EIDE/ATA interface is 16 bits wide.
- Offers excellent performance at a low cost. Top of the line EIDE/ATA drivers that approach SCSI performance levels are only now being developed.
- Very limited device attachment.
- Cannot handle scatter operations well, which is fairly important in OSs that use Virtual Memory operating systems.
- Know the difference between ATA and ATAPI. ATAPI is an extension of the ATA protocol. ATAPI is designed to let your computer recognize CD-ROMs and Tape Drives as though they were additional, ordinary hard disks.

What is a Jumper

- A jumper is a small, usually plastic tap that connects a pair of prongs. These prongs are electrical contact points set into the computer motherboard or an adapter card.
- When you set a jumper, you place a shunt on the prongs that completes a contact. Also called "Shunting" or Jumpering".

Jumper settings: (Master and Slave)

- Each IDE channel can support either one or two devices.
- With IDE devices you need a way of differentiating between devices on a specific channel.
- You do this by designating devices to be either a master or a slave, and then the controller is responsible for addressing commands and data to either the master or the slave.
- The main difference between master and slave drives is that the PC will assign drive letters to the master drive before the slave drive. Your master may be designated a C:\ while the Slave may be designated as a D:\. Just remember that the master gets a drive letter first.
- Each drive's jumper setting can be different, so make sure you look at the drive's documentation to find the proper jumper settings (usually located right on the drive itself either on top or right next to the Jumpering pins on the top, bottom or side of the drive)
- Making both drives the master, or both the slave can generate problems (like not booting)
- The cable makes the decision when you are using Cable Select (CSEL)

Cable Select (CSEL)

- The CSEL system will allow the cable connectors to determine which device is master and which is slave.
- The intention is to save the end user from having to change jumpers when changing a hard disk from master or slave.
- Both devices should have a CSEL Jumpering position and they both have to be configured for CSEL.
- Use a special 3 connector cable (one is connected to the motherboard)
- The other 2 connectors are set so that one of them will tell a drive to be the master and the other will tell its drive to be a slave.
- The drives can be switched by simply changing the connector on the cable they are using.
- The biggest problem with CSEL is that it is non-standard and few PCs use it.
- Another problem is that the cable is much harder to find since it is not often used.
- The IDE cable is relatively short and generally inflexible.

Standard	Details
ATA **(AKA ATA 1, "Plain" ATA)**	Short for AT Attachment, it is a disk drive implementation that integrates the controller on the disk drive itself. Also known as IDE, it supports one or two hard drives, a 16-bit interface and PIO modes 0, 1 and 2. Does not support ATAPI, non-hard-disk IDE/ATA devices, block mode transfers, logical block addressing, Ultra DMA modes or other advanced features and other enhancements. This standard is old and obsolete, so much so that drives developed to meet this standard are no longer made. On the recommendation of the T13 Technical Committee, ATA-1 was withdrawn as an official ANSI standard in 1999.
ATA 2	Supports faster PIO modes 3 and 4, multiword DMA modes 1 and 2, logical block addressing (LBA) and block transfers. Marketed as Fast ATA and Enhanced IDE (EIDE). This standard also is obsolete.
ATA 3	This was a minor revision to ATA 2. ATA-3 improves the reliability of the higher-speed transfer modes, which is an issue because of the low-performance cable used in EIDE/ATA. ATA 3 introduced Self-Monitoring Analysis and Reporting Technology SMART. ATA-3 is not the same as Ultra ATA to ATA 33.
ATA 33 (aka Ultra-DMA and DMA-33)	Supports multiword DMA mode 3 running at 33 MBps. Click here.
ATA 66	A new version of ATA that will double ATA's throughput to 66 MBps. Here is Seagate's FAQ on ATA 66.
ATA 100	Maximum data transfer rate of 100 MBps. To use ATA 100, you have to use a special 80pin ribbon cable. It is backward compatible with other ATA standards. Using the 80pin cable on ATA66 and ATA33 will still improve performance. Click here for link to Dell.

The Three Hots (Spare, Plug and Swap) and Fail over

Hot Spare – A drive you have on hand that can be placed in a server. This refers to the drive itself. A hot spare can become a hot swap/plug.

Hot Plug/Swap – Replacing a hard drive, CD-ROM drive, power supply, or other device with a similar device *without* shutting down the server. Hot plugging is supported by Universal Serial Bus (USB), IEEE 1394 and PCMCIA.

Caution: Do not confuse this Hot Plug with Hot Plug PCI, which is the ability to plug a device into a PCI slot while the PCI bus remains online.

Fail Over – When one device, database, server or network fails, a standby takes its place. This is an important fault tolerance function in a mission-critical environment where constant accessibility is a must.

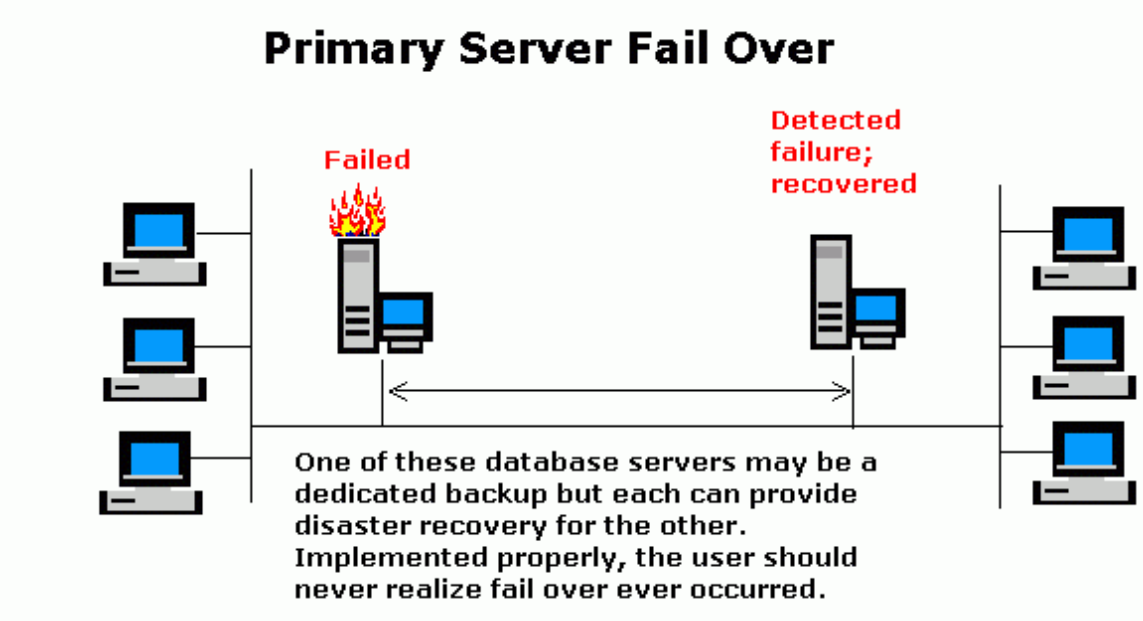

Clustering, Scalability and High Availability

These concepts are not so easily understood as may appear on first glance. You'll notice clustering, scalability and high availability are not separate, distinct concepts. They feed into each other. So do not take this section for granted.

Clustering

A combination of hardware and software solutions used to produce a group of redundant, mirrored servers in a network that will fail over if the primary device fails or is shut down. This is protection against various types of disasters.
Use for Load Balancing and high availability, especially to achieve 99.999 availability (called the five nines; see below.).
An optimally designed cluster should appear, to the user, to be a single system.

Scalability

If your cluster is designed properly, it will be able to change and grow. This ability to continue to function beyond its original context is called scalability.
Scalability upward, rather than downward/backward is the norm. Scaling downward usually means trying to achieve the same goals in a more constrained environment, which is not a good practice.
Note: On the Server+ test, if something, anything, is not a good practice, then it's a wrong answer.

High Availability

The percentage of time a server or cluster is available, generally measured per year, when it falls within an acceptable range (otherwise the availability would not be high).
Optimally, the percentage should be between 99.9% and 100 percent (note: 100 percent is not practical but should still be a goal). The most often mentioned percentage, in IT, is 99.999 percent, AKA, the "five nines."
Five nines works out to 5.39 minutes of total downtime - planned or unplanned - in a given year.
Issues that affect the amount of downtime include:

- **Continuous Availability:** The customer expects 24x365 non-stop service with no interruptions. This is the ideal state because high availability does not imply continuous availability.
- **Fault Tolerance**: Striving for continuous availability, even when hardware and/or software fails. Characterized by redundancy in hardware, including CPU, memory, and I/0 subsystems. This is another ideal state because high availability does not imply fault tolerance.
- **Single Point of Failure (SPOF):** When the loss of a single hardware or software component results in the loss of service. These are components that are not backed up by redundant components and should be.
- **Fail over:** Eliminating SPOF with redundant components that will take over should the primary component fail.

Three ways to achieve the five nines

- Establish special systems designed for high availability. Components within the system would be duplicated for fail over. However, this can be expensive as it would require redundant components that seldom are used
- Establishing mirrored system, so if one fails, the other takes over.
- Then there's clustering, in which matched components do not necessarily have to be duplicates of each other. One of the suggested ways of linking a cluster is through Fibre Channel (see below).

For excellent articles on the five 9s, click here and here.
For the IEEE Task Force on Cluster Computing, click here.
For a free Web-based training class on High availability, click here.
For basic troubleshooting advice, click here.
For an article that explains why you shouldn't be afraid of the command prompt, click here.
For the hows and whys on the distribution of resources on a network, click here.

Fibre Channel

- Fibre Channel is a gigabit interconnect communications protocol designed to meet the ever increasing demand for high performance information transfer. It allows multiple communications between workstations, mainframes, server, data storage and other devices, all using either IP or SCSI technology. Fibre Channel tries to combine the benefits of channel and network technologies. The goals of Fibre Channel include:
- Allowing many well-known existing channel and networking protocols to run over the same physical interface and media
- High bandwidth (100MB/s and beyond)
- Flexible topologies
- Connectivity over several kilometers
- Support for multiple data rates, media types, and connectors

Here is Fibre Channel, bit by bit.

- Fibre Channel is a serial data transfer architecture.
- Designated ports log in to each other through the "fabric," which can be a circuit switch, an intelligent active hub or a loop.
- The Fibre Channel Arbitrated Loop (FC-AL) is the most prominent Fibre Channel standard. FC-AL was designed for new mass storage devices that have high bandwidth and supports full-duplex data-transfer rates of 100MBps. FC-AL uses optical fiber to connect devices and is compatible with SCSI, though it is expected to replace SCSI on high-performance systems.

- Fibre Channel can transmit data between devices at a rate up to 1 Gbps.
- Devices can be ten kilometers (six miles) apart on optical fiber.
- Fibre Channel does work with coaxial and twisted pair cabling but the drawbacks are obvious.
- Fibre Channel offers 3 interfaces: point-to-point, arbitrated loop and cross-point or fabric switched.

Point-to-Point

- Simplest of all topologies.
- A bi-directional link that connects nodes.
- Usually will underutilize the communication link's bandwidth.

Arbitrated Loop

- A form of fabric topology.
- Offers shared bandwidth, which means more than two nodes are connected.
- Fairness algorithm is used to guarantee access to the loop.
- If any link in the loop fails, communication between all L Ports is terminated.
- Less expensive than cross-point.

Cross-point or Fabric Switched

- Highest grade of performance and connectivity. Connectivity is guaranteed with no congestion.
- Bi-directional connection between a node (N_port) and fabric (F_port).
- Multiple paths between F_ports can be configured as non-blocking.
- Efficiently shares bandwidth.
- Point-to-point channel bandwidth is not reduced, even when stations are added to the fabric.
- General Fabric requirements are defined by the standard FC-FG.
- Switch Fabric requirements are defined by the standard FC-SW.

Here is an excellent article on Fibre Channel.
Fibre Channel Industry Association's home page.

AFT - Adaptive Fault Tolerance

- AFT will provide automatic adapter redundancy.
- When the primary adapter goes down, the secondary takes over.
- Adapter Fault Tolerance supports up to four adapter teams, with two to four adapters per team.
- AFT provides an additional safety backup link between your server and your hub or switch.
- If a failure occurs in a hub, switch port, cable, or adapter, uninterrupted network performance will be maintained.
- AFT has a primary adapter and a backup/secondary adapter.
- In normal operation, the backup has "transmit" disabled.
- To use AFT, your adapters must be installed in a Windows NT 4.0 or NetWare 4.1x server-based operating environment.

For a fuller explanation of Adaptive Fault Tolerance, click here.

ALB - Adapter Load Balancing

- Adaptive Load Balancing technology offers an easier, better way to move much more data between a server and the network.
- ALB increases server bandwidth up to 800MBps.
- It is implemented by installing a team of adapters in the server.
- The multiple adapters provide automatic emergency-backup links to the network.
- If a link goes down due to anything from a broken cable, bad switch port, or a adapter failure, the other adapters will instantly pick up the additional load.
- Server operations are not interrupted, and an alert is generated to let the staff know there is a problem.
- Adaptive Load Balancing supports up to four adapter teams (two to four adapters per team).
- Works with any 100BASE-TX switch.
- Adapter teams configured for ALB also provide the benefits of AFT (Adaptive Fault Tolerance).
- Receive rates will remain at about 100 Mbps.

Describe Adapter Teaming

- Teaming is when you install two or more NICs in a server and logically group them so that they will appear to your OS as one network interface.
- This has many benefits including, extra network bandwidth and the fact that you have one MAC address instead of multiple MAC addresses.

- You can also assign a 'virtual' IP address to the team of adapters and not have to rely on complex subnetting.
- Fail over protection with an adapter teaming solution is also available.
- This way your network can withstand having an adapter failure, which gives you "high availability".
- An advantage of adapter teaming is that you can "team" 10/100 Mbps NICs. (This gives you a choice if you want to upgrade to Gigabit Ethernet.)
- Remember: Teaming and load balancing NICs can offer two major benefits: fault tolerance and increased bandwidth.
- In Windows NT, teaming options cannot be implemented on adapters that have been configured for VLANs.
- NetWare can support teaming options and VLANs on the same adapters.
- Adapter Teaming options require NT 4.0 with Service Pack 3.0 and the NDIS driver hotfix from Microsoft.
- Currently, adapters can't support Teaming and Priority Packet configurations at the same time.

For Microsoft's take on Adapter Teaming, click here.

Installations and Upgrades

Getting ready

Know how to install hardware using ESD best practices.
ESD (Electrostatic discharge) is an issue far too many server administrators ignore. There are established best practices, which are:

1. Always work on a server from a grounded, protective rubber mat. If you don't have such a mat, use an antistatic wrist-strap to ground yourself to the server chassis.
2. ***NEVER*** carry cards, memory or anything with a circuit board across a carpeted floor in your bare hands. First, place them in an antistatic bag, and then walk all you want.
3. Use tools designed not to generate ESD, such as screwdrivers with handles treated with rubber.
4. If the server is plugged in – and especially if it's in operation – leave the cover in place. Always shut down the server and unplug it before removing its cover. Don't think this is important? Well, consider this: ATX power supplies keep the motherboard supplied with a constant, but low, voltage, even when the power is off.
5. No eating or drinking around a server.

Cramsession article about ESD.

HCL: Hardware compatibility List

An HCL is a list maintained and regularly updated by hardware and software manufacturers of all hardware known to be compatible with a given operating system.
Before implementing a solution, you should check all available HCLs to identify and prevent incompatibility issues before they occur.
Here are a few examples of HCLs.
Microsoft's HCL
Novell's HCL
IBM's HCL

BIOS upgrade

Obviously, when you first install the system BIOS, you should look for and install any and all available BIOS upgrades. However, after it's been up and running a time, the server's BIOS can become dated. This can become especially critical when it's time to add processors. You'll find the BIOS upgrade in two places:

1. Manufacturer's website (usually the latest available).
2. Installation CD (usually at least a version or two behind updates available online)

PC Guide's take on System BIOS and Advanced Chipset Features. Please carefully read the warning in the second article.

When BIOS/firmware upgrades should be performed

- Always ensure you are at the proper and most up-to-date level of BIOS revision. (Y2k, etc.)
- Do not upgrade a BIOS without first researching it on the vendor's web site, as well as your OS version web site. The reason behind this is that some revisions are not compatible with certain OSs and hardware.
- Always make sure you have the correct revision before flashing (most flash packs are smart enough not to let you do an improper upgrade).
- ***NEVER*** interrupt the flashing process, as it may leave your BIOS or motherboard inoperable and it may need replacement.
- Normally an upgrade should be done when you are upgrading hardware. Most often, newer hardware needs special upgrades in the firmware. (ACPI 2.0)
- To find out what revision of BIOS you have, record it as the server boots, go into the BIOS itself or look through the operating system (MSINFO, etc.) to see what level you are on. Then go to the vendor's site and find out what you need.

- If it is a mission critical server, it is wise to have a standby server with a full backup ready to go. A failed upgrade of your firmware may leave the server down until a replacement for the hardware is put in.
- Generally you would make a set of bootable floppies by downloading a package from the vendor's web site. This floppy or set of floppies will allow you to flash your BIOS and upgrade it.
- Each vendor usually has a site with a set of instructions on what to install and how.

See the Cramsession article on flashing the BIOS.

ACPI

- ACPI (Advanced Configuration and Power Interface)
- To get this support on older servers, you may need to upgrade the firmware to a new revision that has ACPI support. (Go to your vendor's web site support page for reference.)
- ACPI is an industry specification for the efficient handling of power in computer systems and each device's system resources, including IRQs and I/O addresses.
- ACPI also specifies how the BIOS, OS and your peripherals communicate power usage and support with each other.
- ACPI must be supported by the BIOS, motherboard and the operating system.
- With ACPI you can select between power schemes. This allows the user to control the power to specific devices, such as the monitor or Hard drives.
- Your BIOS **must** be ACPI compliant and the operating system you are using must be ACPI-compatible for you to receive ACPI functionality.
- ACPI is designed to work with Windows 98 and with Windows 2000.
- The user can specify how long a device (monitor/hard disk) should stay on.
- The OS has the ability to lower the clock speed when your applications don't need full processor clock speed.
- The OS has the ability to reduce motherboard and peripheral power.
- The computer has a "stand-by" mode, but will leave modem-power enabled to receive incoming faxes.
- Plug-and-play devices are allowed by ACPI, and as soon as they're plugged in and configured, they too can be controlled by ACPI.

Adding a CPU

Installation and Motherboard Compatibility Issues

- Verify that the board is compatible with the CPU.
- Visit the motherboard / CPU vendor website for updated information on their products and any new compatibility issues.
- Verify that the power supply is correct for the motherboard and is also compatible with the CPU.
- If you are adding several items to your system all at once, such as a new processor, more memory, and a new hard disk- do not put them in all at once. Add one item at a time and then test the system between each newly added item. (This will narrow down your troubleshooting if a problem occurs)
- Verify, via the vendor's website, that the system memory is tested and approved for the motherboard by the manufacturer of the board.
- Make sure the correct BIOS version for the motherboard is installed. (Check the vendor website for updates.)
- There are potential risks involved in updating a motherboard BIOS: you could render your system inoperable if you do this incorrectly.
- Make sure the CPU fan / heatsink is installed and operating correctly. You NEED a functional fan and/ or heatsink. If you do not have a fan or a heatsink to dissipate the heat, you are sure to have MANY problems with your CPU. Check with the vendor to see what is required.
- You will need a layer of thermal compound between the heatsink and processor.
- Ensure proper airflow through the case.
- Plug up the slots in the back of the case where expansion cards go whenever a card isn't present. This actually helps maintain proper airflow.
- Beware of cheap power supplies. Some of these actually aim the exhaust fan into the case, rather than out.
- A system that locks up periodically could be indicating a thermal problem.
- Always tie down loose cables and ribbons to avoid problems later.
- For systems that appear to be having an issue with the CPU, install the processor in another compatible motherboard in a known working system. This will help you to isolate an issue with the CPU or motherboard. (ISOLATE!)
- Many new motherboards can configure the processor settings automatically (e.g., VID –- voltage identification -- is one such feature).
- Some motherboards may have you configure the processor settings by jumpers. (Sometimes the settings are configured through the BIOS instead.)
- ALWAYS refer to your motherboard documentation for all the correct jumper settings for the processor installed.

- Research System bus frequency (generally either 66MHz, 100MHz, or 133MHz) and the bus frequency multiplier (the correct multiple of the system bus to achieve the correct processor frequency). This is set in the BIOS itself or with jumpers. Check your documentation to find the correct settings. If this is set incorrectly, you could OVERCLOCK your processor.

For a very good and in-depth review of CPUs: click here, here and here.
Cramsession article on overclocking the processor.

Multiprocessing functionality

Multiprocessing is coordinated processing by two or more computer processors.

- Two types of multiprocessing: SMP (Symmetric Multiprocessing) and MPP (Massively Parallel Processing).

SMP

- Processors share both memory and the I/O bus or the data path.
- A single copy of the OS is in charge of all the processors installed.
- Does not usually exceed 16 processors (unless the OS is ported to higher standards. Some OEM implementations of Windows 2000 Datacenter Server can support 32 CPUs).

MPP

- Up to 200 or more processors can work on the same application.
- Each processor has its own OS and memory, which is why it's also called the "shared nothing system."
- An "interconnect" arrangement of the data paths will allow the messages to be sent between CPUs.
- Setup is rather complicated.

Caution: Multiprocessing should not be confused with multiprogramming.
See MIT's Alewife multiprocessing project.

Cache memory and how it works in a server

Processor Cache

- Processor cache memory is a small amount of high-speed SRAM that will greatly improve your CPU performance.
- This cache resides between the CPU and the main system memory.
- There are two levels of processor cache: "Primary" and "Secondary."
    - Primary cache, (Level 1 or L1) is built into the CPU itself and is usually very small.
    - Secondary cache, (Level 2 or L2) is external from the CPU and usually operates at 12ns to 25ns.
- Motherboard designers usually make this secondary cache upgradeable.
- Cache rarely uses parity, and additional chips called TAG RAM are used.
- "TAG RAM" or "Dirty Bit" RAM will help the cache controller to keep track of where the data is located (within the cache). It is also responsible for the delaying of the transfer of data from cache on its way back to your main RAM.

Write-back Cache

- Caching in which changes made to cached data are not made to the original data at the same time.
- Write-back is also called a copy-back cache.
- Changes to data stored in the L1 cache aren't copied to main memory until absolutely necessary.
- Available on many microprocessors, including all Intel processors in the 80486 family.
- Offers somewhat better performance than write-through because it reduces the number of writes to main memory, but with a slight risk that data may be lost in a system crash.

Write-through Cache

- Caching in which changes made to cached data are made to the original data at the same time.
- Data is written to main memory and the L1 cache at the same time.

Checkout two excellent Cramsession articles on cache, here and here.

Types of memory and server memory requirements

Memory Interleaving	The process of organizing main memory into two or more sections. This way, the CPU can access different sections at once and not wait for memory to catch up through its wait states. Interleaving memory is a way to make up for the slower speeds of DRAM
ECC	ECC (either "error correction or correcting code" or "error checking and correcting"). ECC means that data that is being read or transmitted will be checked for errors and, if necessary, corrected immediately. It differs from parity checking in that the errors are detected and also corrected. ECC is being worked into the design of data storage and transmission hardware because data rates are increasing. As data rates increase, the amount of errors will therefore increase as well.
EDO	Extended Data Output RAM. EDO RAM is a type of DRAM that improves memory performance by eliminating wait states. EDO RAM was initially optimized for the 66 MHz Pentium.
Unbuffered vs. buffered vs. Registered	**Buffered:** when there is so much memory the chipset needs assistance to deal with the large loading introduced by the large amounts of memory. A buffer isolates the memory from the controller to minimize the load the chipset sees. **Unbuffered:** where the chipset controller deals directly with the memory. There is nothing between the chipset and the memory as they communicate. **Registered:** delays memory information for one clock cycle. This is to ensure all communication from the chipset is collected by the clock edge This provides a controlled delay on heavily loaded memories.

ECC with Data Storage

- When a unit of data (or "word") is sent to and stored in RAM, a code that will describe the bit sequence in that word is calculated and then stored along with that unit of data.
- Seven extra bits are needed for each 64-binary digit word.
- When that word is requested, a code is again calculated using the original algorithm.
- The newly generated code is compared with the code generated when that original word was first stored. If both codes match, then this means that the data is free of errors and is sent.
- If the codes do not match, then the missing bits are determined through the comparing of both codes and the data will be supplied or corrected.
- There is no attempt to correct the data that is located in storage.
- If you have reoccurring errors in the same location, an error is sent indicating that you may have a hardware issue or problem.
- ECC increases the reliability without adding a lot of cost.
- Reed-Solomon codes are commonly implemented; they're able to detect and restore "erased" bits as well as incorrect bits.

System Bus Architectures: the basic attributes, purpose, function and performance

Bus Architecture	Details
PCI Bus Mastering	Bus Mastering is usually built into a separate I/O controller chip. Data is sent directly between a controller card and the other devices on the Bus (bypassing the CPU), which greatly reduces the CPU processing. Takes some of the load off of the CPU and improves performance.

Hierarchical PCI Bus	Hierarchical PCI bus has many I/O advantages. A two-level hierarchical bus consists of a primary I/O bus connected to a secondary, I/O bus. The primary is generally HIGH performance while the secondary is of LOWER performance. The primary I/O bus is 32-bit PCI.
Peer PCI Bus	Some systems have 2 PCI buses. Devices on both buses must appear in the IRQ Table.
I2O –Intelligent Input-Output **I2O Org**	Also called Intelligent I/O. This is a new I/O architecture in development by a consortium called the I2O special Interest Group (SIG). I2O is designed to eliminate I/O bottlenecks, by using special I/O processors or IOPs. IOPs handle the many details of interrupt handling, data transfers and buffering. I2O is designed to work with PCI.
Hot Plug/Hot Swap	The ability to add and/or remove devices to a system while the system is up and running. The operating system should automatically recognize a change. This capability allows server administrators to isolate, remove and re-insert a new PCI card without ever shutting down the system. Hot-pluggable or swappable PCI enhance the standard PCI architecture with hot-plug capabilities.
PCI Expansion Slots	Most PCI systems support 3 or 4 PCI slots. Offers a greater variety of expansion cards compared to older slot types. The most commonly found cards are video cards, SCSI host adapters, and NIC cards. Certain functions cannot be provided on the PCI bus, (like serial and parallel ports). They remain on the ISA bus. This is, for the

	most part, because serial and parallel ports are too slow to keep up with the PCI bus.
PCI Interrupts	A PCI bus will use its own internal interrupt system to handle requests from the other cards located on the bus. These interrupts are usually called "#A", "#B", "#C" and "#D" (this is done to avoid confusion and mix up with the other system IRQs), but they are sometimes called "#1" - "#4". If other cards need these interrupts, then they are mapped to regular interrupts (usually IRQ9 - IRQ12). In systems that have more than four PCI slots, two or more of the PCI devices will share an IRQ.
EISA ISA	EISA stands for Extended Industry Standard Architecture. ISA cards will work in EISA slots and not the other way around. Has a 32 Bit Bus Width. The EISA bus supports bus-mastering adapters for greater efficiency.
AGP Overview by Intel Tutorial by Intel	AGP (Accelerated Graphics Port) is a type of interface that enables 3-D graphics to display very quickly. AGP uses your computer's main memory (RAM) to refresh what the monitor image supports: texture mapping, z buffering, and alpha blending required for normal 3-D image displays. AGP main memory use is dynamic. When RAM is not being used for AGP's accelerated graphics functionality, RAM is given back to the OS or other applications and functions.

See the Cramsession article on the PC bus.

How the System Bus affects the performance of a server

- If your bus does not have enough room and I/O traffic is very high (which it usually is on a production server), the system bus can create a bottleneck.
- Always plan your server hardware implementations accordingly.

Management Protocols: basic attributes, purpose, and function.

SNMP

- Simple Network Management Protocol.
- A set of protocols for managing complex networks.
- Works by sending messages, called Protocol Data Units (PDUs), to different parts of a network.
- SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.

For SNMP a' la Microsoft, click here.
For Linuxport's take on SNMP, including a breakdown of the five PDUs, click here.

DMI

- Desktop Management Interface.
- DMI was created by the Desktop Management Task Force (DMTF).
- DMI is used to automate system management and is particularly beneficial in a network-computing environment where there are many computers.
- DMI is hardware and operating system-independent.
- DMI is independent of a specific management protocol such as SNMP.
- DMI is easy for vendors to adopt.
- DMI is mappable to existing management protocols such as the SNMP.
- DMI is used on network and non-network computers.

For SUN tutorial on DMI, click here.
For IBM Tutorial on DMI, click here.
For a comparison of SNMP and DMI with Microsoft's Windows Management Instrumentation, click here.

Backup, restore and disaster recovery

Disaster is an occurrence causing widespread destruction and distress; a catastrophe or a grave misfortune.
Source: dictionary.com.

Guidelines

- A disaster can be as catastrophic as a tornado destroying the primary operation site or misfortunate as the loss of critical data. How a network administrator will cope with an inevitable disaster will depend on the disaster and the amount of planning for disaster recovery.
- Disaster recovery refers to the process of restoring a system after a "disaster".
- Sometimes a disaster recovery is a simple as restoring a backup.
- Other disasters require planning such as having standbys and spares of all critical pieces of replaceable hardware components for a potential server failure in stock and tested.
- Always have a documented disaster recovery plan.
- Always keep a set of the data OFFSITE.
- Establish an alternative site (cold site; offsite new building - maybe even a different state)
- Prepare a special grouping of people to work at your alternative site and devise a plan to get them to the new site from the disaster site.
- Decide what products are needed to support the recovery process.
- Some disaster recovery scenarios even include having a complete duplicate of your server standing by, in case of disaster. (Cold or Hot standby)
- Simple items such as a UPS can save you a lot of headaches, such as in the case of power failures.
- Disaster-recovery management should begin with a planning meeting.
- Ensure that the procedures are current and accurate.
- Update your documentation and test it regularly.
- Consider a rotating schedule of different technical staff members.

See the Cramsession article on how to test and replace batteries in the APC UPS 1000.

Questions to ask when developing a disaster recovery plan

- What will the company need if disaster strikes?
- What department(s) has priority in getting back online first?
- How much of the data is at risk?
- What is the minimum and maximum downtime you can afford?
- What is your cost per minute if your system is down?

- Are there redundant networks that will replace your failed system?
- Would a clustered environment minimize your risk?

For Windows 2000 Server disaster-recovery concepts, click here.

Backup strategies

Full	Backs up all data. Takes the longest to backup, but fastest to restore. Clears archive bit on all backed-up files.
Incremental	Only backs up files added or changed since last backup. Quick to back up, but restore requires multiple tapes (last full backup tape plus all incremental tapes since the last full backup. When restoring incremental, always take care to restore the correct order, oldest to newest.). Clears archive bit on all backed-up files.
Differential	Only backs up files since last FULL backup. Each backup takes a little longer, but restore requires only two tapes (most recent full backup and most recent differential backup). Does not clear archive bit.

NEVER, NEVER, NEVER MIX INCREMENTAL AND DIFFERENTIAL.

GFS (Grandfather, Father, Son) backup rotation

- Most common of all backup methods.
- Requires a minimum of 21 tapes (assuming a 5-day rotation; 4 daily, 5 weekly, and 12 monthly.) for a full year of backups. This number excludes spare tapes for the inevitable failure and the cleaning tape.
- Create a monthly backup (Grandfather). This is to be stored permanently offsite, never to be reused.
- Create a full weekly backup (Father)
- Create a differential or incremental daily backup (Son)
- Daily and weekly backups are reused each week and can be stored on or offsite.

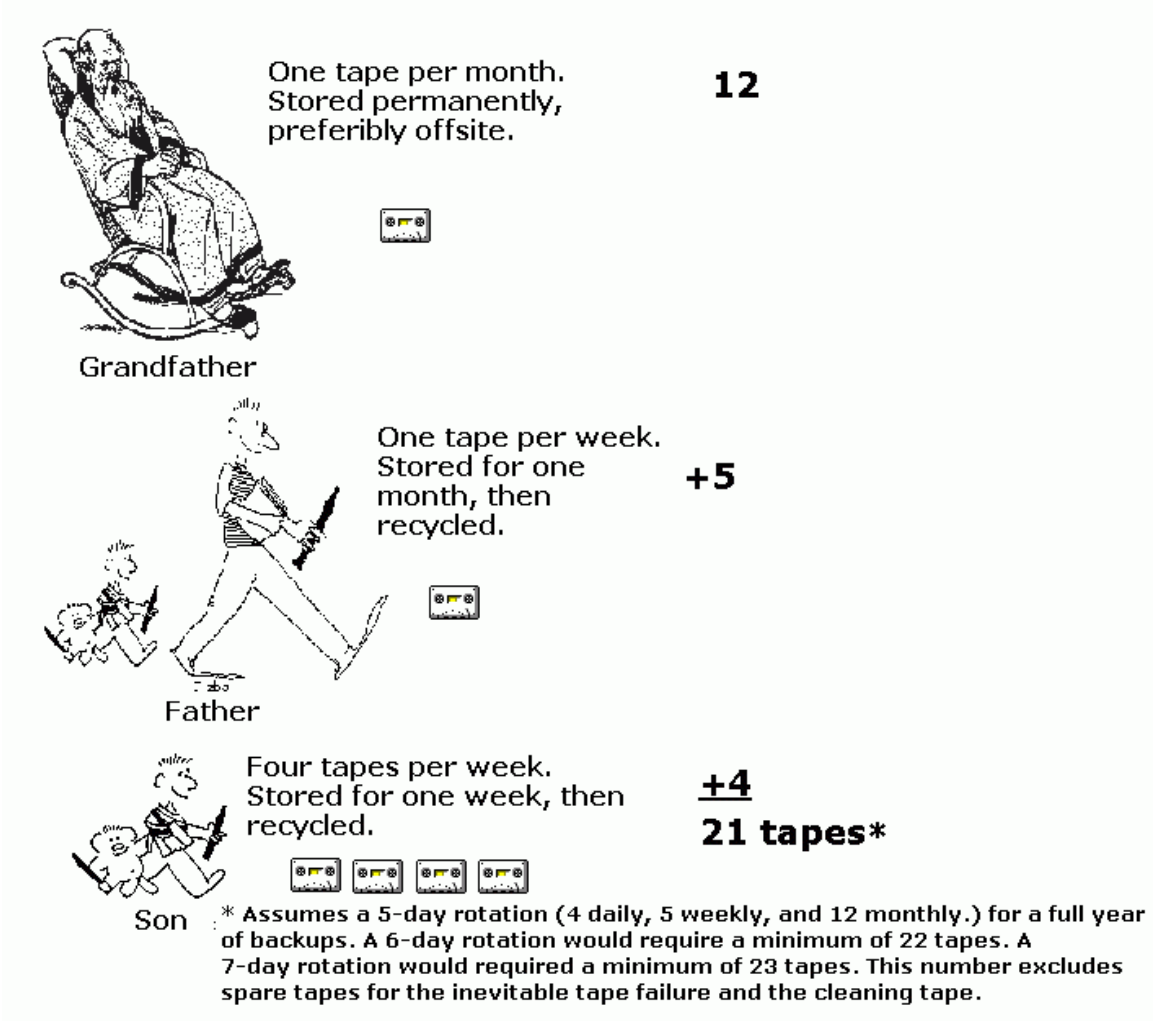

Bear in mind that how many tapes a site/client needs will depend on the number of days in that client's backup rotation. For instance, if the client wants a quarterly rotation (based on a four-week month, four daily tapes (son), four weekly tapes (father) and four monthly tapes (grandfather) for a total of 16 tapes for a full year of backups) then that's what that client requires. Under the original GFS rotation, only three tapes were used (Day 1 (son), Day 2 (father) and Day 3 (grandfather), then the cycle would repeat). Obviously, this type of rotation is not very practical for long-term data recovery but, if that's what the client wants, then that's what that client requires.

Backup Drives

Quarter Inch Cartridge – Introduced in 1972, primarily for the UNIX market, QICs are cheap – cheapest of all listed here - and excellent for stand alone PCs but not the best choice for networks. They are slow and capacity is limited. The name is derived from the width of the tape – quarter inch. The drive comes in two sizes: 1) 3.5-inch mini cartridge that holds up to 20 GB, and 2) a 5.25-inch data cartridge that stores up to 50 GB.

Digital Audio Tape (DAT) – Standard for companies that require high-speed, reliable and simple storage devices. This drive uses the same technology you'll find in your VCR. Comes in two formats:

1. **Digital Data Storage (DDS):** The current DDS standard, DDS-4, uses a 120-meter cartridge tape that stores up to 40 GB.
2. **Digital Linear Tape (DLT) –** High-capacity tape storage that stores up to 80 GB on a single tape.

Please note: DDS and DLT are often confused as competing storage formats. While DLT has greater capacity than DDS, they are both DAT storage media.

Travan – an enhancement of QIC, newer network storage versions hold up to 10 GB of compressed data.

Mammoth – Introduced in 1996 by Exabyte, three drives have a 20-GB capacity with a transfer rate of 3 MBps. Mammoth-2 tapes, introduced in 1999, offer a new technology called SmartClean. When the drive is dirty – a common problem with tape drives – the drive will clean itself without user intervention.

How the capacities compare

Storage Medium	Capacity
Diskette	1.44 MB
iomega Zip	100 MB or 250 MB
CD-RW	700 MB (MB for MB, most expensive here)
iomega Jaz	1 GB or 2 GB
DDS-1	2 GB
DDS-2	8 GB
Travan	10 GB
DDS-3	12 GB
QIC 3.5-inch mini cartridge	20 GB
Mammoth	20 GB
DDS-4	40 GB
QIC 5.25inch data cartridge	50 GB
DLT	80 GB

For an excellent Intel article on GFS, click here.
Here's a Cramsession article on why backing up data is so important and another article on the Tao of backup.

Backup storage and restore

Hot and cold sites

Hot Site – the primary operation site; this is where the network actually is on a daily basis.
Cold Site – a backup location, preferably with redundant equipment and resources, which can become hot should the primary hot site become unavailable (as in a disaster).
Benefits

- Allows work to continue, even in a disaster, thus still striving for the 99.999 uptime goal.
- Provides development site available for testing or for some shared critical resources.

Disadvantages

- It's expensive.
- Cold sites typically are designed only for temporary work.

Planning for redundancy
When ordering servers with an eye towards redundancy, there are many key areas to consider:

- Hard drives
- Power supplies
- Cooling fans
- NICs
- Multiple processors
- UPSs
- Generators
- Cluster servers

Server room Security: concepts and the importance of physically securing a server

Physically Securing a Server

- There are many ways to secure a server.
- Lock the console of the operating system you are using. (password protected)
- Lock the cabinet that the server is mounted or sitting in.
- Lock the server room (preferably with digital locks).
- Limit access to the server room to authorized server administrators only.
- Remember, the greatest danger to the server room comes not from outside the company but from within. While there are more external threats, the attacker who most likely can breach your security is an internal security risk.

General Security Concepts

(These can apply to both network and server security.)

- A strong commitment from management to provide sufficient resources to get security established and to enforce security policies and procedures is absolutely necessary.
- A staff of personnel trained and dedicated to security tasks is needed.
- A well-defined security mission statement that EVERYONE will read.
- A well-developed "security awareness" training program.
- Clearly defined, implemented and documented security policies and procedures, which are supplied to everyone within your organization.
- A strong flow of information to and from the appropriate groups.
- A security incident response team.
- External and internal security perimeter controls (e.g., firewalls)
- A suite of host- and network-based security auditing and improvement tools.

Questions to ask about server security

- Are you concerned about insider threats?
- Should you trust your users? What level of computer knowledge do they possess?
- Are most of your users local or remote? How do they access the servers?
- What kind of permissions and access control do you need to set on your servers?
- Are all backdoors disabled? (THERE ARE MANY!)
- Are any bugs or flaws in system software enabling your users to find holes in your security? (Get on these types of lists! NTBUGTRAQ, NT Security)
- Are you testing your own security policies (proactively attempting to get into the same servers that you are charged with protecting)?
- How much sensitive information is on the servers? What is the loss to your organization if this information is compromised or stolen?
- Do you need different levels of security for different parts of your organization?
- Are there security guidelines, regulations, or laws your organization is required to meet?
- Do business requirements take precedence over security where there is a conflict? (If they do, this is not good).

See the Cramsession article on network security.
See this Cramsession article on how to ensure your server gets hacked.

Server Room Environment

You'll need to know the difference between server room security and server room environment. In general, discussions concerning the server room environment include:

Temperature

- Servers produce an amazing amount of heat and it's important to make sure that heat is dissipated quickly and efficiently so the servers don't overheat.
- An ideal temperature is high 60s, low 70s (Fahrenheit), about 20 degrees Celsius.
- Keep fans on hand, in case of emergency (such as your cooling system breaking down). Removing ceiling tiles in such a situation can help – a little – till the HVAC folks arrive and put your server room back to right.

Humidity

- Too much humidity is bad for anything metal, especially when it holds a circuit board. On the other hand, you should never let the server room get "too dry."
- Install a humidistat and dehumidifier.

ESD

- Elevated floors are common in many server rooms because they provide ESD protection through the flooring framework. That framework is usually grounded out with stout copper wire going to the heavy earth ground. Thus, it's not hard to see how this provides a stable ESD platform to work from.
- Workbenches should be protected with grounded rubber mats.
- Use wrist straps when working inside a server, especially with any equipment that has a circuit board on it.

Power Surges

- Any good motherboard is just one electrical spike away from trash.
- Equip the server room with surge protectors.

Back-up Generators

- These are good to have in the event of a local power loss but are considered too expensive in smaller operation centers.
- To get an idea of whether a generator is a good strategy, compare the cost of purchase and implementation with that of how much the company would lose if it lost power for a 24-hour period.

Fire suppression

- Sprinklers are a no-no in server rooms. Sure, the water will put out the fire but not without destroying the heart of your network; then there's the ever present danger it will leak and destroy part of your network without a fire at all.
- Then there's Halon, which is very good at suppressing a fire in the server room without destroying valuable equipment, but is deadly to humans. The server room techs have only seconds to find the exits before they die. Halon also is ozone layer depleting.
- An alternative to Halon is FM200 (heptaflouoropropane, AKA HFC-27ea). For a good description of this product, click here.

Other concerns

- Critters in the server room, such as insects, rats and mice, are environmental concerns.
- Due to flooding hazards, server rooms never should be located in a basement or on the first floors of buildings near rivers, lakes and other bodies of water.
- However, in earthquake zones, placing the server room on an upper floor is a no-no. If you're in an earthquake zone that happens to be in a flood prone area, then priorities must be set accordingly.

Maintenance Issues

Describe Shutdown Procedures Across the Following Network Operating Systems

Operating System	Shutdown Procedures:
Novell NetWare (Broadcast a message)	**At the console:** Down (Typing "down" will ensure system and data integrity before turning off file server power. In Novell speak, this is how you "elegantly" shut down the server). Restart server
Microsoft Windows NT/2000	Go to the start button (Or Ctrl alt del) and select Shutdown.
UNIX/Linux UNIX	There are various ways to shutdown a Unix / Linux server. This one is pretty basic: Unix SCO ACE: Shutdown now -g –0
IBM OS/2 Stupid OS/2 tricks	Shutdown OS/2 from a command line Just type **SHUTDOWN** from any OS/2 command window. If that doesn't work, the command may not be on your path. Try **X:\OS2\INSTALL\SHUTDOWN**, where X is your main OS/2 drive.

Shutdown basics

- Remember to save all your work and close all running applications.
- Broadcast a message to clients that they should save their work and give a time when the server is going down.
- You can (with the right permissions) also remotely shutdown your servers manually, by application or by script.

Rack mounting servers

KVM implementation APC's KVM Raritan's KVM pictures	Keyboard / Video / Mouse Preplan your KVM setup before running cables and mounting servers. You should label what goes where and get a general idea of the layout. Cabinets get tight so allocate proper space. Also, most KVM mount towards the top of the cabinet.
Cable management Racks: APC's Racks and Hardware	Cables get messy when dealing with racks and cabinets. Make sure you run them along special rack ducts and label them. This would aid you if you blew out a rack mountable monitor and had to get in and put a "spare" setup in. Properly labeling these cables will allow you to route your cables easily into the backup hardware. If you don't, then you will be fishing for them. LABEL CABLES
Proper layout and Security Lockable cases / racks	LOCK THE SERVER UP! An inaccessible server is a fairly protected one. Always lock the console / secure the console when not in use. Lock the rack door and monitor who has access to the server or where it is located. You will have enough to worry about with network and O/S security to deal with. Lock it up.

Measuring in "U"

- A "U" is the open, cutout side of a rack-mounted appliance where the screw fits in to hold the appliance to the rack. Racks generally have an interior dimension of about 19 inches. What the average tech is interested in is the usable space or U.
- A U = 1.75 inches.
- A mini cabinet holds about 23U; medium cabinet about 31U and tall cabinet about 40U.
- Servers generally listed at 1U, 2U or 3U. A server with a 5.25-inch hard drive may list at 4U. This means it will take up four holes, or Us, or about 7 inches.

The importance and use of maintenance logs and service logs (documentation) Importance of Logs:

- Logs contain information the Operating System and its applications put on file (usually a text file or a *.log).
- They store CRITICAL information that you can (and will) use to solve moderate to very serious server related issues.
- Almost all products 'log' their behavior. When you install software, it logs. If you boot your server, it may log. You can use these logs to troubleshoot.
- Other types of logs you can keep are records that you as a server administrator will keep to record all maintenance that is done on the server's O/S or hardware. This is highly important (and useful) for many reasons.
    - An example: Servers are usually not powered down (you would schedule an outage), but you need to put in a hot fix or a critical update immediately for your NIC drivers. What NIC is it? What are the specs? If you kept logs on all this, you would have an answer.
    - Another example: After a weekend off you come to work. The server crashed for an unknown reason. You desperately try to resolve the issue but where do you start? Well, look in the Maintenance and service logs to see if any work was done on the weekend during an outage. This may help fix your problem.

Major Features of Server Operating Systems

Server Operating Systems:	Major Features of latest releases:
Novell NetWare	Netware 6.x – newly released version Netware 5.x Netware 4.x Netware 3.x
Microsoft Windows NT/2000	Windows NT Windows 2000
UNIX/Linux	Red Hat SCO ACE
IBM OS/2	OS/2 Warp Server

Also, check out these Cramsession OS Crams. I kid you not. Know what's on these Crams as well as this one. Server+ contains questions relevant to each of these operating systems.
Windows 2000 Pro
Windows NT

NetWare 5.1
RedHat Linux
Unix
IBM OS/2

To help you see the differences, check out this Microsoft VS Novell white paper
Other Server Platforms: FreeBSD | MacOSX | OpenLinux | Sun Solaris | Full List

Define the Impact of Domains and Directory Services on Server Hardware Design

- When designing your server or server-based environment, you must plan out what hardware requirements you will need in advance.
- If you were planning to install a PDC (Primary Domain Controller) for Windows NT, you would assume that the one thousand users in house would be requesting authentication from it. This would require higher hardware requirements.
- If you were planning a database (Oracle / SQL) or some kind of "Directory Services" server, you would need higher hardware requirements.
- Always take all hardware into account when planning, but always pay close attention to the memory and disk space requirements.
- Always follow your documentation on what is needed for the O/S or applications. (Some platforms, like SQL and PROXY, can be VERY memory intensive.) Domain controllers, DNS servers and Directory Services servers should always have a large amount of disk space and memory.

Network Interface Issues

Describe the following common network interface protocols:

LAN Protocol Type:	Details:
Ethernet 10Base2, 10Base5, 10BaseT	A Network access layer or a physical layer protocol. Ethernet networks use CSMA/CD. Ethernet will run over a variety of cable types at 10 Mbps. Ethernet is very similar to IEEE 802.3 standards.
Fast Ethernet 100BaseT, 100BaseT4, 100BaseTX	FE offers speed increases, to 100 MBps, over 10BaseTEthernet. FE will keep the frame formatting, MAC mechanism, and MTU intact. Based on an extension to the IEEE 802.3 specification.
Gigabit Ethernet 1000baseT / 802.3ab DELL Gigabit Ethernet solution 3Com's White paper	1000BaseT runs over single and multi-mode copper fiber. Supports distances of up to 100 meters with four pair CAT-5 copper cable. Will auto-negotiate between 100MBps and 1GBps speed
Token Ring	Developed and supported by IBM. Token Ring runs at 4 or 16 Mbps on a ring topology. Token ring is Similar to IEEE 802.5 standards.

Check out Cisco's Documentation on ETHERNET and TOKEN RING.

NICs as they relate to the system bus

- Very old systems use ISA but the NIC is still plugged into a PCI slot.
- In older types of systems, your PCI bus that the NIC card inserted into was a significant I/O bottleneck for bandwidth and significantly reduced efficiency and speed.
- Now, in newer systems (Gigabit Ethernet), that bottleneck is reduced and the bandwidth and efficiency is greatly increased.

For 3Com's viewpoint on NICs vs. the System Bus, click here.

Basic Troubleshooting

CONFIG	Novell - IP Stack configuration viewing
IFCONFIG	Unix - IP Stack configuration viewing
WINIPFCG	WIN95/98 - IP Stack configuration viewing
IPCONFIG	NT/2000 - IP Stack configuration viewing
PING	End to End connectivity test - ICMP ECHO
TRACERT	Used to find bottleneck - How many hops
NETSTAT	Used to see listening ports
ARP	Used to view current ARP Cache

For troubleshooting basic TCP/IP problems, click here.

Special thanks to
Robert J. Shimonski for the original material

and to Karen Ann Robertson-Kidd
for major revisions to this Cramsession.